I Z A Institute of Labor Economics

Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 17263

Medical Residency Subsidies and Physician Shortages

Cici McNamara
Mayra Pineda-Torres

AUGUST 2024

I Z A Institute
of Labor Economics

Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 17263

Medical Residency Subsidies and Physician Shortages

Cici McNamara
Georgia Institute of Technology, Atlanta

Mayra Pineda-Torres
IZA and Georgia Institute of Technology, Atlanta

AUGUST 2024

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

Medical Residency Subsidies and Physician Shortages*

We quantify the impact of federal subsidies for graduate medical education on primary care physician (PCP) supply by examining the impact of Section 5503 of the Affordable Care Act, which increased the number of residents that teaching hospitals in rural and high-need areas could receive subsidies for training. Instrumenting for selection into the program using its eligibility and allocation criteria, we find that the provision increased both the recruitment of residents into primary care and time spent at teaching hospitals in high-need areas, resulting in an increase in PCP supply in treated counties of 5.2 percent.

JEL Classification: I18, I28, J24

Keywords: Medicare, Affordable Care Act, primary care

Corresponding author:
Mayra Pineda-Torres
School of Economics
Georgia Institute of Technology
221 Bobby Dodd Way NW
Suite 220, Atlanta
Georgia 30332
USA

E-mail: mpinedat@gatech.edu

* We thank Carson Cole and Tehreem Hussain for their excellent research assistance. We also thank Adibah Abdulhadi, Eli Liebman, Jason Lindo, Sara Markowitz, Kelly Noonan, Michael Richards, and Tingting Zhang for their useful comments. All errors are our own.

1 Introduction

Primary care physician shortages are a current and growing concern in the U.S. (AAMC, 2021a). Baby boomers aging into Medicare and the Affordable Care Act (ACA)'s expansion of health insurance coverage have both contributed to increased demand for primary care over the past decade (HHS, 2022; SSA, 2007). However, the supply of primary care physicians (PCPs) has not grown at the same rate as demand, and all primary care specialties are projected to be in shortage by 2035 (HRSA, 2022). Physician shortages in rural areas are of particular policy concern. In 2020, the Health Resources and Services Administration estimated that 83 million Americans live in an area with a shortage of PCPs and that 62 percent of these Health Professional Shortage Areas (HPSAs) were rural (GAO, 2021). Figure 1 illustrates this point by plotting the trend in all physicians and PCPs per capita separately for rural and urban counties. Physician density is higher in urban counties every year. Furthermore, the growth rate for all physicians is higher in urban counties than in rural ones, while primary care density is actually decreasing in rural counties. These patterns are concerning, as PCP density has been repeatedly linked to lower morbidity, lower mortality, and longer life expectancy (Basu et al., 2019, 2021; Ku and Druss, 2020; Macinko et al., 2007; Shi et al., 1999; Starfield, 1991).

Figure 1: Trend in average physician supply in rural and urban counties

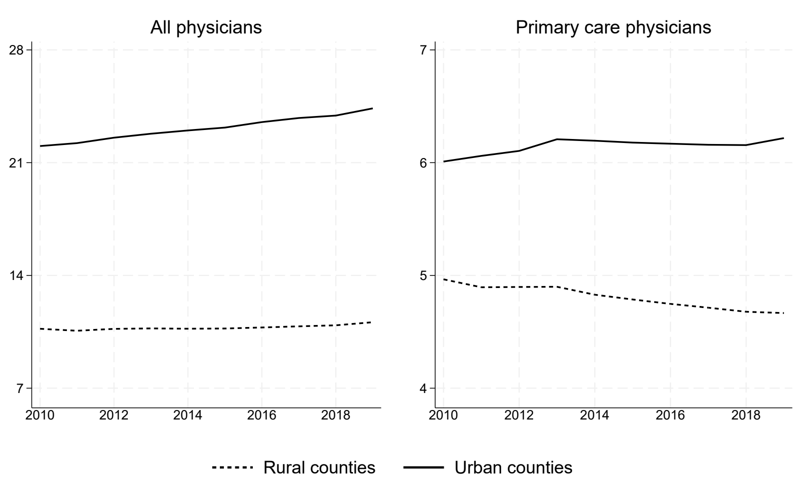

Source: Authors' analysis of Area Health Resource Files. *Notes:* Plotted is the population weighted average of physicians per ten thousand people across counties. Primary care physicians include those specializing in family medicine, general practice, internal medicine, and pediatrics.

One way to address the PCP shortage is to train more physicians in primary care. Physicians typically select into a specialty at the end of medical school and spend their graduate medical education (GME), more commonly known as "residency," training in that specialty. The federal government subsidizes this training through Medicare payments. Revising the formula for GME subsidies is one way by which the federal government may influence the specialties and locations in which residents train. Given that 57% of medical residents who completed residency training between 2011 and 2020 now practice in the state where they completed their training, there is scope for policies aimed at increasing resident training in underserved areas to have a long-run impact on physician supply (AAMC, 2021b).

In this paper, we quantify the impact of an increase in GME subsidies on the level and geographical distribution of PCP supply by examining the impact of Section 5503 of the ACA. Section 5503 aimed to increase the number of primary care residents trained in high-need areas by reallocating residency subsidies away from teaching hospitals in areas with relatively high physician supply to those in areas where supply was low, which included but was not limited to rural areas. We instrument for hospitals' selection into the reallocation program using an instrumental variables model that relies on the published criteria used by the Centers for Medicare and Medicaid Services (CMS) to determine which hospitals were eligible to apply for a subsidy increase and the order in which applications were considered. We separately estimate models at the program, hospital, and county levels to quantify the policy's effects on recruitment into primary care, where residency training takes place, and where attending PCPs choose to practice.

We find that in each year following the implementation of Section 5503, treated hospitals increased their teaching intensity by an average of 3.91 full-time equivalent residents per 100 beds, a 9 percent increase over the baseline average. This increase is in part attributable to an estimated 23 percent increase in the number of residents recruited into primary care residency programs located in areas that were prioritized by the policy's allocation criteria. We then examine whether the policy was successful at impacting PCPs' long-term location decisions and increasing attending PCP supply in underserved areas. Using the same empirical model employed in our hospital- and program-level analyses, we show that Section 5503 resulted in an increase in county-level attending PCP supply of 5.2 percent. These estimates imply a "conversion rate" of 31 percent, by which we mean that for every primary care resident recruited as a result of the Section 5503, there is a corresponding increase in medium-run attending PCP supply of 0.31.

This paper contributes to our knowledge of how provider payments affect healthcare provider input choice. Studies have found a positive relationship between Medicaid payments to nursing homes and both staff size and quality measures (Konetzka et al., 2004; White, 2005; He et al., 2020), as well as between Medicaid payments and both staff numbers and hours (Cohen and Spector, 1996; Feng et al., 2008; Lin, 2014; Foster and Lee, 2015). Kaestner and Guardado (2008) exploit changes in Medicare reimbursement for nurses generated by geographic reclassification of hospitals to quantify the effect of subsidies on nurse utilization and patient outcomes and finds no meaningful effect on either. In a paper closely related to ours, Nicholson and Song (2001) find that the revision of the GME payment policy to compensate teaching hospitals for the indirect costs of training residents resulted in an increase in teaching hospitals' resident utilization.

This paper also contributes to our understanding of the determinants of physicians' specialty and location choices. Falcettoni (2017) studies the effect of loan forgiveness and salary incentives on physicians' location choices taking specialty as given, finding that salary incentives are relatively more effective at increasing the share of physicians that choose to practice in rural areas and that residents strongly prefer to practice close to where they completed their residency training. Kulka and McWeeny (2019) estimate that the implementation of loan forgiveness programs increase the number of physicians in rural counties by three. They also find that loan forgiveness incentives do not succeed in getting physicians to move out of their home states, suggesting that these programs may be limited in their ability to address physician shortages in rural areas. Other studies have documented peer effects from medical school classmates (Arcidiacono and Nicholson, 2005) and economic factors including indebtedness, expected earnings, income, expected relative hours worked, and length of the training period (Bazzoli, 1985; McKay, 1990; Thornton, 2000; Nicholson, 2002; Thornton and Esposto, 2003) as determinants of specialty choice. Heterogeneity in the preferences of medical school graduates for specialty characteristics generates the potential for regulations to influence the composition of specialties' residents. For instance, (Wasserman, 2023) finds that the implementation of a 2003 cap on weekly hours worked by residents resulted in an increase in entry by women into specialties with binding average weekly hours.

In examining medium-run effects of Section 5503 on attending PCP supply in the medium run, this paper also contributes to our knowledge of work location choice that has estimated home bias in employment decisions (Greenwood, 1997; Diamond, 2016). The relevant location bias in this paper is for the area where a physician completed their residency training. Fadlon et al. (2020) consider a similar type of location bias, showing that female physicians' first job assignment affects their location decisions in the long run in the context of Denmark. Our results indicate a conversion rate for residents recruited under Section 5503 of 31 percent. The conversion rate we consider here may be the result of the retention of residents recruited under Section 5503 in the areas in which they trained or the reallocation of attending PCPs to treated areas to serve as complements to newly recruited residents in the production of hospital services. While we are unable to disentangle the relative contribution of these two mechanisms to our estimated conversion rate, the fact that 31 percent serves as an upper bound on the retention mechanism suggests that this type of location bias is modest. Lastly, this paper contributes to a large literature on the effects of the Affordable Care Act and, in particular, to the relatively small subset of this literature that studies the supply-side effects of the healthcare reform.

Our estimates of the efficacy of GME subsidies at increasing PCP supply in high-need areas are important for deciding how to prioritize this policy in relation to others, which include loan forgiveness, the scope of practice laws, and telehealth (Falcettoni, 2017; Kulka and McWeeny, 2019; Markowitz and Adams, 2022; Panzirer, 2021). They are also important given that GME subsidies remain a popular choice among this set of policy alternatives. For example, beginning in 2023, the Consolidated Appropriations Act started increasing the subsidies of selected teaching hospitals by up to 200 cumulative slots using eligibility criteria similar to those used under Section 5503.[1] We also note that healthcare is not the only labor market in which the government intervenes to expand supply: state and federal governments subsidize teaching certification and have subsidized private firms' training of new workers in the manufacturing sector (Holzer et al., 1993; Georgia Student Finance Commission, 2023). Our results therefore speak more broadly to the efficacy of subsidizing the cost of training borne by firms in addressing labor shortages.

Our paper is structured as follows. Section 2 provides information on federal funding for GME, Section 5503, and the resident labor market. Section 3 describes our data sources. Section 4 discusses our empirical model and identification. We present results in Section 5 and conduct a counterfactual analysis in Section 6. We conclude in Section 7.

[1] Public Law 116-260

2 Background

2.1 Federal funding for graduate medical education

Medical school graduates are required to undertake at least three years of GME before practicing medicine independently as attending physicians. The federal government subsidizes GME through Medicare payments to teaching hospitals. These payments are the largest source of government support for residency programs and totaled about $15 billion in 2018. Medicare GME payments compensate teaching hospitals for both the direct and indirect costs of operating residency programs. Direct graduate medical education (DGME) costs include resident salaries and benefits, accreditation and licensing fees, and faculty compensation. Indirect medical education (IME) payments are meant to compensate teaching hospitals for the fact that residents are less efficient at providing care. The average GME payment made by Medicare to a teaching hospital in 2018 was $11 million, $3.3 million of which was for DGME costs and $8.7 million of which was for IME costs (GAO, 2017).

Teaching hospitals receive retrospective DGME and IME payments based on the number of full-time equivalent (FTE) residents that they trained in the previous year.[2] DGME and IME payments are increasing in the number of FTEs a hospital trains up to a cap, beyond which the marginal effect of training an additional resident on reimbursements is zero.[3] These caps were established by the Balanced Budget Act of 1997 and for most hospitals equal the number of FTEs that the hospital was training in 1996.[4] For example, the total payment to be made by Medicare to reimburse a teaching hospital for the direct cost of graduate medical education in year t is

$$\text{DGME payment } t = \min\begin{Bmatrix}\text{FTEs} & \text{FTEs} \\ t & 1996\end{Bmatrix} \times \begin{pmatrix}\text{Per-resident} \\ \text{amount}\end{pmatrix} \times \begin{pmatrix}\text{Medicare} \\ \text{patient load } t\end{pmatrix}$$

[2]A resident will typically count as one FTE during their actual residency and as less than one during their fellowship (MedPAC, 2001).

[3]More precisely, DGME and IME payments are both increasing functions of the hospital's three-year FTE rolling average, where any given year's FTE count cannot exceed the hospital's cap. The parameters of the payment formulas depend on the hospitals' number of beds, direct costs of residency program operation in 1984, and Medicare's share of total inpatient days (CMS, 2018).The DGME payment formula is DGME payment $=$ (Three-year FTE rolling average) $\times$ (Per-resident amount) $\times$ $(\frac{\text{Part A inpatient days}}{\text{Total inpatient days}}) \times (\frac{\text{Part C inpatient days}}{\text{Total inpatient days}} \times 0.86)$.
The IME payment formula is IME payment $= (1.35 \times ((1 + \frac{\text{Three-year FTE rolling average}}{\text{Number of beds}})^{0.405} - 1) \times$ (DRG payment).

[4]Public Law 105-33

where FTEs t and FTEs 1996 are the number of FTE residents that the hospital trained in years t and in 1996, respectively; the per-resident amount is a time-invariant hospital-level multiplier; and the hospital's Medicare patient load is Medicare's share of inpatient days. We henceforth refer to a hospital's FTE count as of 1996 as its "resident cap." In 2018, 70 percent of teaching hospitals were over either their DGME or IME cap, meaning that they trained more residents than they were receiving funding for through Medicare (GAO, 2017). Figure 2 shows the average number of residents trained by each hospital in a year and the number of residents for which it was reimbursed.[5] The average number of residents trained is consistently above the average cap, which remains largely flat over the sample period.

Figure 2: Divergence between resident counts and caps

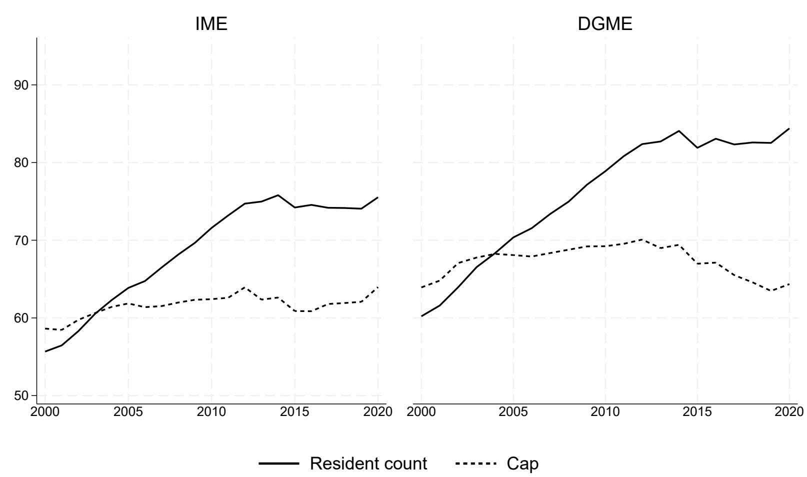

Source: Authors' analysis of CMS Healthcare Cost Report Information Systems hospital cost reports. *Notes:* The solid black line labeled "Resident count" shows the average number of FTE residents trained, where the average is computed across hospitals within a year. The dashed black line shows the average number of FTE residents the hospital could receive a GME subsidy from Medicare for training. IME stands for Indirect Graduate Medical Education. DGME stands for Direct Graduate Medical Education.

2.2 Section 5503 resident cap redistribution

In July 2011, CMS implemented a redistribution of subsidized resident slots from hospitals that had been consistently operating below their caps to hospitals in eligible areas.[6] Hospi-

[5] DGME and IME FTE counts for a given hospital may not be equal, as not all types of time enter into the computation of both counts. For example, time spent doing research may be counted differently for the computation of DGME FTEs than for IME FTEs; see 42 CFR. §412.105(f)(1)(iii)(B).

[6] Public Law 111-148, §5503

tals operating below their caps had 65 percent of their excess residency slots revoked. 267 hospitals saw a cumulative reduction in their IME and DGME caps of 628.05 and 726.08, respectively. This pool of revoked slots is small in comparison to the number of available residency program slots in any given year, which is over 20 thousand (NRMP, 2022). It is large, however, in comparison to the average baseline resident cap and number of residents trained by a hospital in the baseline period, which were approximately 70 and 80, respectively.

Hospitals could apply for a cap increase from the pool of revoked slots if they satisfied certain eligibility criteria. Seventy percent of the revoked slots were made available to hospitals located in states in the bottom quartile of the resident-to-population ratio distribution, while the remaining 30 percent of slots went to hospitals located in rural counties and those in the top ten states in terms of primary care Health Professional Shortage Area (HPSA)-to-population ratios.[7] Panel A of Figure 3 plots states in ratio space. The ratio thresholds for eligibility are in bold on each axis and are illustrated with dashed lines. All hospitals in states to the left of the vertical or above the horizontal dashed lines were eligible to apply for a cap increase along with those in rural counties.

Eligible hospitals that applied for a cap increase were considered in order of their place in the distribution under which they qualified:

> We also proposed that, in determining which applicant hospitals receive priority within the priority category of hospitals located in a State in the lowest quartile for resident-to-population ratios that hospitals in a State that is ranked lower in the quartile (with number one being the lowest) would receive preference over hospitals in states that are still within the quartile, but ranked higher. For example, all other things being equal, a hospital located in Montana would receive preference over a hospital located in Idaho, while this hospital would receive preference over a hospital located in Alaska, and so on. Similarly, we proposed that, in determining which applicant hospitals receive priority within the priority category of hospitals located in a State that is among the top 10 of these areas in terms of the ratio of Primary Care HPSA population to total population, hospitals in an area that is ranked higher in the top 10 (with number 1 being highest and number 10 being lowest) would receive preference over hospitals in an area that is still within the top 10 but ranked lower. For example, all other things

[7] In order to be considered a primary care HPSA, an area's population to provider ratio must be at least 3,500 to 1 (KFF, 2023).

being equal, a hospital located in Louisiana would receive preference over a hospital located in Mississippi, while a hospital in Mississippi would receive preference over a hospital located in Puerto Rico, and so on (Federal Register Vol. 75, No. 226, p. 72181).

Figure 3: Illustration of Section 5503 eligibility and allocation criteria

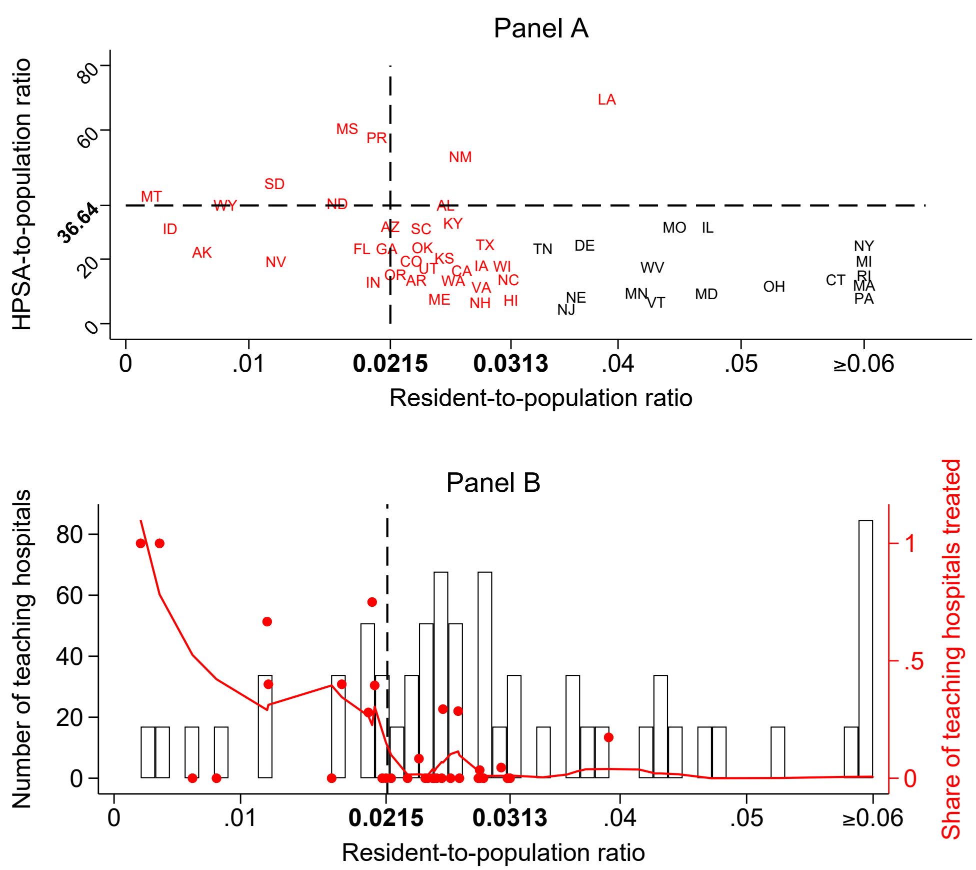

Source: Authors' analysis of CMS Healthcare Cost Report Information Systems hospital cost reports and eligibility and allocation criteria from Federal Register Vol. 75, No. 226. *Notes:* An observation in Panel A is a state. The abbreviations of states included in the analytical sample are written red, while those excluded from the analytical sample are written in black. In Panel B, bars outlined in black show the density of teaching hospitals by the resident-to-population ratio of their state. Red dots show the share of teaching hospitals in each state that receive a cap increase under Section 5503, while the red line provides locally weighted scatterplot smoothing of the share of hospitals treated. Vertical and horizontal dashed lines illustrate the eligibility cutoffs for the resident-to-population and HPSA-to-population ratios, respectively.

58 of 225 eligible hospitals received cap increases after the allocation criteria were applied. Eligible hospitals may not have received a cap increase either because they did not apply for one or because they did apply but all available slots had been allocated prior to the consideration of their application. Panel B of Figure 3 shows the density of teaching hospitals across states according to their resident-to-population ratio as well as the share of teaching hospitals in each state that received a cap increase under Section 5503. As described in the quote above, the share of eligible hospitals that received a cap increase falls with the resident-to-population ratio. For example, hospitals in Georgia that applied for cap increases received none because the entirety of the 70 percent pool made available to hospitals in low resident-to-population ratio states had been exhausted by the time applications from these hospitals at the top of the bottom quartile of the distribution were considered.[8] Panel B also shows a mass of hospitals with resident-to-population ratios very close to the eligibility threshold of 0.0215. Some of these hospitals are eligible to apply for a cap increase by virtue of their HPSA-to-population ratio, as shown in Panel A, but many are ineligible despite having very similar values for the eligibility criteria as their eligible counterparts. This bunching is a byproduct of eligibility being determined by each state's rank in the ratio distribution rather than the value of the ratio itself. We will leverage these features of the eligibility and allocation criteria in the construction of our control group and instrument.

Hospitals that received cap increases were required to use at least 75 percent of their increase to fund new primary care or general surgery residency positions. Hospitals whose caps were raised under the reallocation could not solely engage in "cap relief" and use the increase to fund already-filled slots. CMS would audit treated hospitals after five years to ensure that the 75 percent threshold was being met and could revoke cap increases if they were not. This policy was therefore "slot-neutral" in that it did not change the total number of residents that Medicare could potentially subsidize, but it was not budget-neutral in that it increased the total number of filled subsidized slots. Despite this 75 percent requirement, it is still empirically ambiguous whether treated hospitals would be able to increase their program size for many reasons, the most salient of which is the potentially inelastic supply of residents to rural areas and primary care relative to urban areas and other specialties.

Figure 4 shows the trend in the resident cap over time for hospitals stratified by whether

[8] For more information see: https://www.cms.gov/medicare/medicare-fee-for-service-payment/acuteinpatientpps/downloads/section_5503_cap_decreases_and_increases.zip

they received a cap increase under Section 5503. The average cap increase was 13 slots, which constitutes about 20% of the baseline cap. Hospitals that received cap reductions under Section 5503 were already operating below their caps and so would have seen no change to their realized reimbursements after the reductions, which were small in size. Indeed, McNamara and Hussain (2023) show that there is no evidence that hospitals receiving a cap decrease under Section 5503 experienced any change in resident utilization as a result of the policy. For this reason, our focus in this paper will be on estimating the causal effect of Section 5503's cap increases, and we exclude hospitals that received cap decreases from all analyses below.

Figure 4: Trend in resident caps for hospitals by treatment status

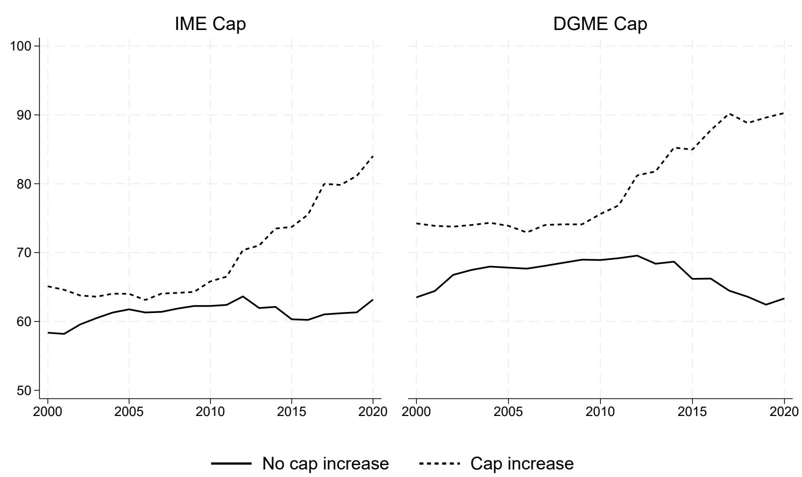

Source: Authors' analysis of CMS Healthcare Cost Report Information Systems hospital cost reports. *Notes:* The solid black line shows the average number of FTE residents for which teaching hospitals that were not subject to the cap reallocation received a GME subsidy from Medicare for training. The dashed black line shows the analogous average across teaching hospitals with a cap increase under the reallocation. IME stands for Indirect Graduate Medical Education. DGME stands for Direct Graduate Medical Education.

While Section 5503's objective was to increase PCP training and supply in high-need areas, there are at least two channels through which spillovers to non-primary care physicians might manifest. First, if PCPs and non-primary care physicians are complements in the production of medical services, then the subsidizing of primary care residents may incentivize the recruitment of non-primary care physicians. Second, Section 5503 stipulated that primary care residents must fill at least 75 percent of additional subsidized slots. The remaining 25 percent may have been filled by non-primary care residents. We consider this potential for spillovers to non-primary care physicians in the construction of our outcome measures below.

2.3 The Match

Medical school graduates are matched to residency programs through a stable matching algorithm administered by The National Residency Matching Program known as The Match. In every year since the 1970s, the number of medical school applicants for positions in residency programs has exceeded the number of available residency slots. In 2010, for example, 28.8 percent of all applicants for residency programs went unmatched (NRMP, 2010). There is variation in the fill rate for residency slots across specialties, especially as it relates to the number of first-year residency slots that are filled by seniors from U.S. medical schools. Figure 5 summarizes the results of the 2010 Match for the twenty largest specialties in terms of number of programs, where programs are sorted in descending order of the match rate for U.S. seniors. Family medicine is the specialty with the lowest overall and U.S. senior match rates, though its overall match rate is still above 90 percent. Internal medicine, another primary care specialty, also has a relatively low U.S. senior match rate of 55 percent.

Figure 5: Match rates across specialties

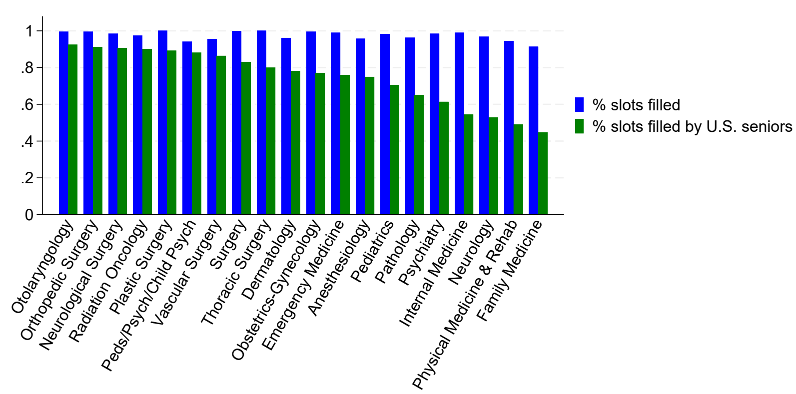

Source: Author's calculations from National Residency Matching Program data on the 2010 Match.

This figure highlights two features of the labor market for residents at the time of the passage of the ACA. First, residency slots were in short supply. Second, primary care specialties were not in high demand by U.S. seniors, and residency slots in these programs were largely filled by other applicants. These facts suggest that resident supply is likely elastic with respect to the addition of new slots but that the marginal resident who chooses to fill a newly added primary care slot is likely to be an international medical school graduate or a previously unmatched U.S. medical school graduate. There is little evidence that hospitals that received cap increases adjusted their wages to attract more U.S. seniors, as can be seen in Figure A1. In fact, Figure A2 shows that real resident wages have remained remarkably stable over the past two decades, exhibiting much less growth than that of other hospital employees.[9]

3 Data

3.1 Hospital cost reports

We use three data sources, the first of which are the hospital cost reports from the CMS Healthcare Cost Report Information System (HCRIS). The cost reports contain financial data for all Medicare-certified hospitals in the U.S. This includes annual data on resident caps, the number of residents trained, and GME payments. All of these variables are measured separately for direct and indirect GME. The cost reports also identify hospitals that received cap increases under Section 5503 and the size of each hospital's cap change.

All GME data in the cost reports is reported at the hospital level and is not broken out by specialty. However, for DGME payments the count of residents trained is broken out into counts for primary care and non-primary care programs. The cost report's definition of primary care includes traditional primary care specialties - family medicine, general practice, internal medicine, and pediatrics - as well as obstetrics and gynecology (OB/GYN).[10] This definition of primary care does not perfectly coincide with that from Section 5503, which does not include OB/GYN but does include general surgery.[11] To the extent that Section 5503 increased general surgery resident utilization, we can expect to see changes in both the count of primary care and non-primary care FTEs as measured by the cost reports. The

[9] This may be a result of collusive practices among teaching hospitals; see Jung v. Association of American Medical Colleges (2005).

[10] 42 CFR §413.75(b)

[11] 42 CFR §413.79(n)

cost reports also contain information on hospital characteristics, including the total number of discharges, number of Medicare and Medicaid discharges, and number of beds. We use these variables as controls in our empirical specification.

We focus on teaching hospitals, which include both hospitals that sponsor residency programs and those at which residents rotate but do not themselves sponsor programs. In 2010, about 20% of all hospitals in the cost report data were teaching hospitals. We restrict our analytical sample to teaching hospitals in rural counties and the states whose name abbreviations are written in red in Panel A of Figure 3. This group includes eligible hospitals as well as "almost eligible" ones, by which we mean that their state's resident-to-population ratio lies just above the threshold for eligibility. In particular, hospitals in states with resident-to-population ratios below 0.0313 were included in the analytical sample to take advantage of the mass of hospitals located between this value and the eligibility threshold of 0.0215. Our results are robust to alternative cutoffs for inclusion in the analytical sample. This approach allows us to leverage the discontinuous drop in the likelihood of treatment at the eligibility threshold in our identification. We also exclude from our analytical sample hospitals that received a cap decrease under Section 5503. Our control group is therefore comprised of teaching hospitals that did not receive cap increases located in eligible or "almost eligible" states.

Column 1 in Table 1 shows the averages in 2010 of characteristics of the teaching hospitals in our analytical sample. On average, these hospitals trained 69.5 residents, 64 percent of which were in primary care. This average number of residents trained is higher than the average resident cap of 55.6 residents. Across all hospitals, the average number of slots awarded and increase in potential GME subsidies were 1.3 and $100 thousand, respectively.

Table 1: Summary statistics and exogeneity of instrument for hospital analysis data

	(1) Mean (S.D.)	(2) Coefficient (S.E.)	(3) p-value
Resident utilization			
Number of residents	69.5 (124.2)		
Residents per 100 beds	23.7 (34.5)		
Share of residents in primary care	63.6 (35.3)		
Section 5503 cap increase			
Number of slots awarded	1.3 (5.5)		
GME payment awarded (mill.)	0.1 (0.5)		
Eligibility and allocation criteria			
Resident-to-population ratio	0.027 (0.009)		
HPSA-to-population ratio	23.8 (13.1)		
Rural	11.1%		
Hospital characteristics			
Number of beds	249.8 (180.7)	2.48 (1.37)	0.0770
Number of discharges (ten thous.)	1.54 (1.23)	1.36 (0.61)	0.0311
Share of discharges from Medicare	28.1 (14.1)	0.06 (0.10)	0.5540
Share of discharges from Medicaid	17.6 (13.7)	0.01 (0.08)	0.8865
Resident cap	55.6 (104.9)	-0.01 (0.003)	0.0023
Total GME payment (mill.)	5.5 (10.2)	-0.06 (0.20)	0.2001
County characteristics			
Unemployment rate	9.5 (2.5)	-0.19 (0.42)	0.6545
Median household income (ten thous.)	4.9 (1.1)	-7.57 (5.38)	0.1164
Share White	75.3 (16.3)	0.01 (0.13)	0.9594
Share Black	15.7 (15.4)	0.13 (0.17)	0.4353
Share in poverty	16.7 (4.6)	0.05 (0.26)	0.8516
Inpatient days per capita	0.97 (0.68)	0.15 (2.18)	0.9447
Number of hospitals	18.7 (26.3)	0.25 (2.23)	0.9124
Number of hospital beds (thous.)	4.10 (6.09)	0.94 (1.78)	0.5988
Share with health insurance	80.72 (5.86)	-0.48 (0.47)	0.3203
Number of observations	551		
p-value of joint F-test	0.1042		

Source: Authors' analysis of CMS Healthcare Cost Report Information Systems hospital cost reports and Area Health Resource Files. *Notes*: An observation in column (1) is a teaching hospital in 2010. Cells show mean (standard deviation) for continuous variables and percents for binary variables. Reported in column (2) is the coefficient and its standard error from a regression of our instrumental variable on the given variable using the full analytical sample. Number of beds, number of discharges, median household income, county-level number of hospital beds, and county-level number hospitals are log transformed for these regressions. Column (3) reports the p-value for each regression. Standard errors are clustered at the state level. The p-value of the F-test for the joint significance of the estimates from a regression of the instrument on all of the hospital and county characteristics is reported in the bottom row. GME stands for graduate medical education. HPSA stands for health professional shortage area.

3.2 Area Health Resource File

We also rely on county-level data from the Area Health Resource Files (AHRF), which contains counts of primary care physicians aggregated from the 2010-2019 American Medical Association Physician Masterfile. These counts are available for all physicians and are also broken out separately for MDs and DOs. We consider heterogeneity by type of medical education (allopathic versus osteopathic) because DOs have been shown to be more likely to practice in rural and underserved areas (Fordyce et al., 2012). We also use the AHRF's data on county demographics, socio-economic status, and healthcare utilization as controls in our empirical specification. Data on each state's resident-to-population and primary care HPSA-to-population ratios as of 2009 come from the federal regulations themselves.[12]

Column 1 in Table 2 presents averages as of 2010 characteristics of the teaching hospitals in our analytical sample. On average, 222 physicians practice in these counties, of which 203 are MDs and 19 are DOs. The cumulative number of slots and potential GME payments across the teaching hospitals in a county awarded as a result of Section 5503 were 1.3 and $100 thousand, respectively. The averages of the county characteristics presented in the bottom Panel Are similar to those observed in Table 1, with the exception of number of hospitals and beds, which are smaller in Table 2. This difference is reflective of the fact that the data in Table 2 has been collapsed from the hospital to the county level.

3.3 National Residency Matching Program reports

We use annual reports from the National Residency Matching Program (NRMP) on the outcomes of each Match to measure the flow of medical school graduates into residency programs in different specialties and locations (NRMP, 2024). Not all teaching hospitals house residency programs and not all residency programs are sponsored by hospitals. While a focal hospital sponsors most residency programs, programs typically entail clinical rotations at multiple hospitals that may be within or outside of their own system.[13] It is therefore possible that an increase in the number of FTE residents trained at treated teaching hospitals represents a reallocation of resident time away from untreated hospitals in a given residency program and toward treated hospitals in the same program without an increase

[12] Federal Register Vol. 75, No. 226, pp. 72177-72181

[13] For instance, in 2007, Grady Memorial Hospital in Atlanta was Georgia's largest teaching hospital in terms of the number of residents trained despite not sponsoring a residency program (GPBW, 2007).

Table 2: Summary statistics and exogeneity of instrument for county analysis data

	(1) Mean (S.D.)	(2) Coefficient (S.E.)	(3) p-value
Physician counts			
All	221.9 (461.4)		
MDs	202.8 (430.1)		
DOs	19.1 (39.3)		
Section 5503 cap increase			
Cumulative number of slots awarded	1.5 (7.3)		
Cumulative GME payment awarded (mill.)	0.1 (0.5)		
Eligibility and allocation criteria			
Resident-to-population ratio	0.029 (0.013)		
HPSA-to-population ratio	23.6 (13.2)		
Rural	34.4%		
County characteristics			
Unemployment rate	9.3 (2.8)	0.006 (0.51)	0.9905
Median household income (ten thous.)	4.6 (1.1)	-4.06 (5.18)	0.4367
Share White	82.0 (16.0)	-0.20 (0.15)	0.1708
Share Black	11.9 (14.7)	0.17 (0.19)	0.3793
Share in poverty	16.9 (5.9)	0.26 (0.32)	0.4191
Inpatient days per capita	0.92 (0.94)	1.23 (0.89)	0.1694
Number of hospitals	5.1 (8.0)	-2.49 (1.82)	0.1767
Number of hospital beds (thous.)	0.95 (1.80)	-2.66 (2.63)	0.3174
Share with health insurance	81.8 (5.27)	-0.62 (0.41)	0.1346
Number of observations	500		
Joint F-test	0.1693		

Source: Authors' analysis of CMS Healthcare Cost Report Information Systems hospital cost reports and Area Health Resource Files. *Notes*: An observation in column (1) is a county in 2010. Cells show mean (standard deviation) for continuous variables and percents for binary variables. Reported in column (2) is the coefficient and its standard error from a regression of the instrumental variable on the given variable using the full analytical sample sample. Number of beds, number of discharges, median household income, county-level number of hospital beds, and county-level number hospitals are log transformed for these regressions. Column (3) reports the p-value for each regression. Standard errors are clustered at the state level. The p-value of the F-test for the joint significance of the estimates from a regression of the instrument on all of the hospital and county characteristics is reported in the bottom row. GME stands for graduate medical education. HPSA stands for health professional shortage area.

in the total number of primary care residents being trained.

In order to assess the impact of Section 5503 on specialty choice, we construct a panel dataset of the number of residency positions offered and filled - known as the quota and number matched - by each residency program in the U.S. between 2006 and 2019 using the NRMP's publicly available reports on the results of the Match.[14] We are able to determine the specialty and location of each residency program but are not able to match residency programs to hospitals and determine which programs are affiliated with treated hospitals, which limits the analyses we can perform with these data. Nonetheless, these data are valuable in allowing us to determine whether Section 5503 influenced the overall supply of PCPs and corroborate the results of our hospital- and county-level analyses.

Table 3: Summary statistics for program analysis data

	Mean/%	Standard error
Program characteristics		
Primary care	34.6%	
Quota	6.25	(5.78)
Matched	5.98	(5.82)
Eligibility and allocation criteria		
Resident-to-population ratio	0.04	(0.05)
HPSA-to-population ratio	24.6	(13.7)
Rural	9.6%	
Number of observations	1963	

Source: Authors' analysis of National Residency Match Program reports. *Notes*: An observation is a residency program in 2010. Primary care specialties include family medicine, general practice, internal medicine, and pediatrics. Quota is the number of positions offered by the program. Matched refers to the average number of positions filled as part of The Match. HPSA stands for Health Professional Shortage Area.

Table 3 summarizes the characteristics of programs included in the NRMP data in 2010 as well as the Section 5503 eligibility and allocation criteria for each program's location. Nearly two thousand residency programs participated in the Match in 2010, approximately one-third of which were in primary care specialties. The average number of residency positions offered was 6.25, while the average number of positions filled as part of the Match was 5.98. Figure 6 shows the trend in the average residency program size for programs stratified by

[14]The dates of publications for these reports are between 2007 and 2020 but pertain to quotas set in the previous year.

whether they were eligible to apply for a Section 5503 cap increase based on their focal hospital's location. The trends show that the quota and number matched for residency programs in eligible areas are below that of programs in ineligible areas until 2011, when there was a sharp increase in the size of eligible programs. While these differences in these means are not statistically different, Figure 6 provides preliminary evidence of the effects we will estimate more rigorously below.

Figure 6: Mean quota and number matched for programs stratified by eligibility for a cap increase under Section 5503

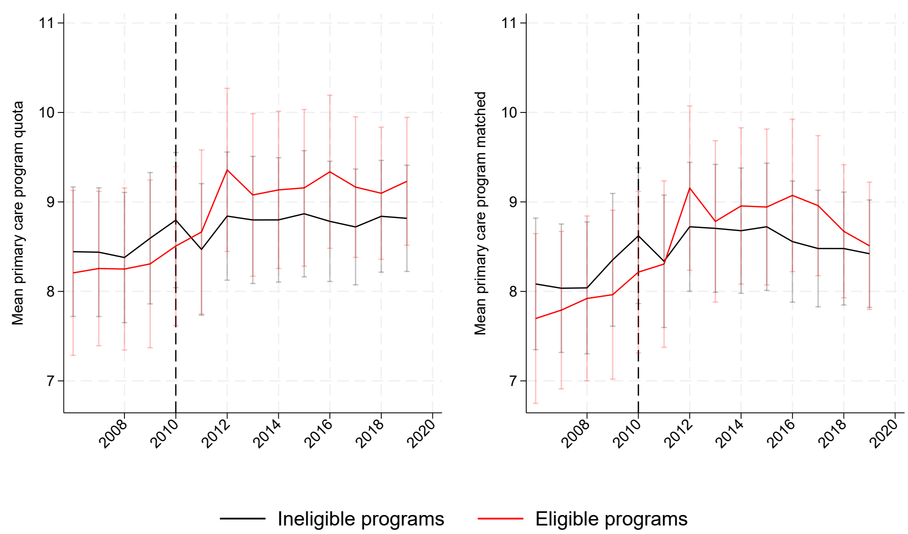

Source: Author's calculations from National Residency Matching Program data. *Notes:* Red lines correspond to the average of the given outcome across residency programs eligible to apply for a residency cap increase under Section 5503, which includes those located in rural counties, those in states in the bottom quartile of resident-to-population ratios, and those in the top ten in state-level HPSA-to-population ratios. Black lines correspond to the average of the given outcome across residency programs ineligible to apply for a residency cap increase under Section 5503. Transparent capped bars provide 95 percent confidence intervals for these means.

4 Estimation and identification

Our estimates of interest are of the effect of the Section 5503 GME subsidy increase on the number and location choices of PCPs in the years following the provision's implementation. We will estimate many of the same models at the hospital, program, and county level, and so we summarize them denoting a unit of observation by the generic subscript i. Assuming Y_{it}

is the appropriate outcome measure for unit i, these effects are given by the λ_t coefficients in the estimating equation

$$Y_{it} = \sum_{\substack{t = t_1 \\ t \neq t_{ref}}}^{2019} \lambda_t^1 \mathbb{1}\{\text{Increase}\}_i \times \mathbb{1}_t + X_{it}'\beta_1 + \eta_t + \eta_i + \omega_{it}^1 \tag{1}$$

where $\mathbb{1}\{\text{Increase}\}_i$ is an indicator for unit i being receiving a cap increase under Section 5503, $\mathbb{1}_t$ is an indicator for the year equalling t, X_{it}' are time-varying controls, and η_t and η_i are year and unit fixed effects, respectively. t_1 and t_{ref} are the initial and reference years for the sample period. Direct estimation of equation 1 is confounded by the fact that hospitals selected into treatment under Section 5503. As discussed in Section 2.2, not all hospitals eligible for a cap increase received one, either because they chose not to apply for this increase or because the reallocated slots had already been exhausted. As a result, hospitals that applied for a cap increase may be different from those that did not apply in ways that are both unobservable and correlated with Y_{it}. For example, hospitals that applied for and received a cap increase may have differentially responded to and participated in ACA provisions other than Section 5503 in ways that affected their teaching intensity, which would bias our estimates of the effect of Section 5503.

We address this endogeneity issue by instrumenting for treatment under Section 5503 using the eligibility and allocation criteria employed by CMS in determining which hospitals could apply for a cap increase and in what order their applications would be considered. Instrumental variables (IV) models have been widely used in other contexts where agents endogenously select into treatment and intent-to-treat estimates may be muted by non-compliance (Angrist et al., 2009; Deming et al., 2014; Angrist et al., 2022; Milligan and Stabile, 2011; Dahl and Lochner, 2012; Currie and Gruber, 1996; Goodman-Bacon, 2021). In particular, we instrument for $\mathbb{1}\{\text{Increase}\}_i$ in equation (1) using

$$Z_{s(i)} = Rank_{s(i)} \times \mathbb{1}\{\text{Eligible}\}_{s(i)}.$$

Here, $Rank_{s(i)}$ is the rank of unit i's state in descending order of resident-to-population ratios. For instance, as can be seen in Panel A of Figure 3, Montana would have the highest value of $Rank_{s(i)}$, Idaho the second highest value, and New York the lowest value. $\mathbb{1}\{\text{Eligible}\}_{s(i)}$ is a binary indicator for being eligible to receive a cap increase by virtue of being in the bottom quartile of state-level resident-to-population ratios (i.e., having a resident-to-population ratio below 0.0215). As can be seen in Panel B of Figure 3, the likelihood that a hospital is treated under Section 5503 is increasing in $Rank_{s(i)}$ and discontinuously falls past the eligibility threshold of 0.0215. Interacting $Rank_{s(i)}$ with $\mathbb{1}\{\text{Eligible}\}_{s(i)}$ captures both of these sources of variation in the likelihood of treatment. Our IV model can be summarized as

$$\mathbb{1}\{\text{Increase}\}_i \times \mathbb{1}_t = \sum_{\substack{t\,=\,t_1 \\ t\,\neq\,t_{ref}}}^{2019} \lambda_t^2 Z_{s(i)} \times \mathbb{1}_s + X_{it}'\beta_2 + \eta_t + \eta_i + \omega_{it}^2 \;\forall t \in [t_1, 2019] - \{t_{ref}\} \tag{2}$$

$$Y_{it} = \sum_{\substack{t\,=\,t_1 \\ t\,\neq\,t_{ref}}}^{2019} \lambda_t^3 \widehat{\mathbb{1}\{\text{Increase}\}_i \times} \mathbb{1}_t + X_{it}'\beta_3 + \eta_t + \eta_i + \omega_{it}^3, \tag{3}$$

where equations (2) and (3) are the first- and second-stage regressions, respectively.

We additionally estimate equation (1) by ordinary least squares as well as the reduced form of the IV model summarized above, which is given by

$$Y_{it} = \sum_{\substack{t\,=\,t_1 \\ t\,\neq\,t_{ref}}}^{2019} \lambda_t^4 Z_{s(i)} \times \mathbb{1}_t + X_{it}'\beta_4 + \eta_t + \eta_i + \omega_{it}^4. \tag{4}$$

In the hospital-level estimation of these models, we use as our outcome Y_{ht} the resident-to-bed ratio of the hospital (multiplied by 100), the primary measure of teaching intensity used by CMS. The numerator of the resident-to-bed ratio may include all residents, primary care residents, or non-primary care residents in different specifications. In program-level analyses, we use as our outcome variable the program's annual quota and number matched. We perform program-level analyses separately for primary and non-primary care specialties. We note that because we cannot map hospitals to programs and determine which residency programs were treated, our program-level analyses are restricted to the estimation of the reduced form equation 4. In county-level analyses, we use the natural log of counts of all physicians, MDs, and DOs separately as our outcomes to estimate the percent change in PCP supply affected by Section 5503. While we use 2010 as the reference year in hospital- and program-level analyses, we use 2013 as the reference year the county-level analyses since increases in area-level attending physician supply can only manifest after the residents recruited in response to the cap increase complete their training. Primary care residencies are typically three years long, so the soonest that effects could manifest is in 2013.

Time-varying hospital and county characteristics included as controls are the share of discharges attributable to Medicare, the share of discharges attributable to Medicaid, the logarithm of total discharges, the fraction of the population that is White, the fraction of the population that is Black, inpatient days per person, percent of the population in poverty, unemployment rate, the logarithm of the total number of hospitals, the logarithm of the median income, and the share of the county's population that is insured. Hospital-level regressions are weighted by the number of beds, while county-level regressions are weighted by population. Standard errors for all models are clustered at the state level.

The validity of this empirical approach relies on two assumptions. First, the eligibility criteria affect the likelihood of receiving a cap increase (i.e., the instruments are relevant), and second, the eligibility criteria affect outcomes only through their effect on the likelihood of receiving a cap increase, conditional on our included instruments (i.e., the instruments satisfy the exclusion restriction). Considering first the strength of our instruments, for each of our first stage regressions, Table B1 provides the cumulative effect of all of the instruments on the given endogenous regressor (i.e., $\sum_{\substack{\mathrm{s=t_1} \\ \mathrm{s} \neq t_{ref}}}^{2019} \hat{\lambda}_s^2$), the standard error of this cumulative effect, and the regression's Sanderson and Windmeijer (SW) F-statistic, which provides a test for weak instruments in models with multiple endogenous variables (Sanderson and Windmeijer, 2016).These results show that a one unit increase in the instrument increases the likelihood of receiving a residency cap increase of approximately 0.08 percentage points for both the hospital- and county-level specifications, and this increase is statistically significant at the 1 percent level. Additionally, all of the SW F-stats are high for all of the endogenous regressors for both the hospital- and county-level specifications, the lowest being 21.6.The Kleibergen-Paap rk LM statistic tests for underidentification and indicates that our model is identified at the five percent level.

We provide further evidence of the exogeneity of our instrument in Tables 1 and 2 by providing the correlation between the instrument and exogenous covariates. These coefficients are all statistically insignificant for the county-level analytical sample and are statistically insignificant for the hospital-level sample with the exception of number of discharges and the size of the resident cap. A joint F test of all of the coefficients leaves us unable to reject the null that they are uncorrelated with the instrument.

5 Results

5.1 Resident utilization

We summarize our main findings by providing event-study plots of the λ_t coefficients from the OLS, IV, and reduced form models presented in Section 4. These plots for our hospital-level analyses are provided in Figure 7, where the outcome variable is residents per 100 beds. In the left column, all residents are included in the numerator of this outcome measure, while in the middle and right columns, only primary care and non-primary care residents are included, respectively. Across all specifications, we cannot reject the null hypothesis that the average of the pre-period coefficients equals zero. The average of the post-period coefficients for the IV models is statistically significant at the 1 percent level and equals 8.35, 3.91, and 4.74 for all, primary care and non-primary care FTEs, respectively. This increase in primary care resident utilization constitutes a 9 percent increase over the baseline average of 42.7 primary care residents per 100 beds at treated hospitals. These effects are larger than those estimated for the OLS model, which equals 3.45 and 2.13 for all primary care residents, respectively. The OLS estimates for non-primary care residents are not statistically significant. Reduced form estimates indicate a one unit increase in the instrument is associated with an average annual increase in resident utilization in the decade following Section 5503's implementation of 0.07 FTE residents per 100 beds, 0.03 of which are attributable to increases in primary care utilization and 0.04 of which are attributable to non-primary care.

In Figures A3 and A4, we present results using as our endogenous treatment variables the size of the cap increase and the natural log of the increase in GME subsidies corresponding to the cap increase. The average of the post-period coefficients for the IV model estimates in Figure A3 show that an increase in a hospital's cap of one slot increases total resident utilization by 0.48 residents per 100 beds and that this increase is generated increases in primary and non-primary care utilization of 0.22 and 0.28 residents per 100 beds. Assuming hospital beds are fixed over the sample period, the results in Panel (a) of Figure A3 indicate that hospitals were, on average, just shy of being in compliance with the requirement that 75 percent of cap increases be used to increase program size by the end of the sample period.

Figure 7: Estimation results for the effect of Section 5503 on teaching hospital residents per 100 beds

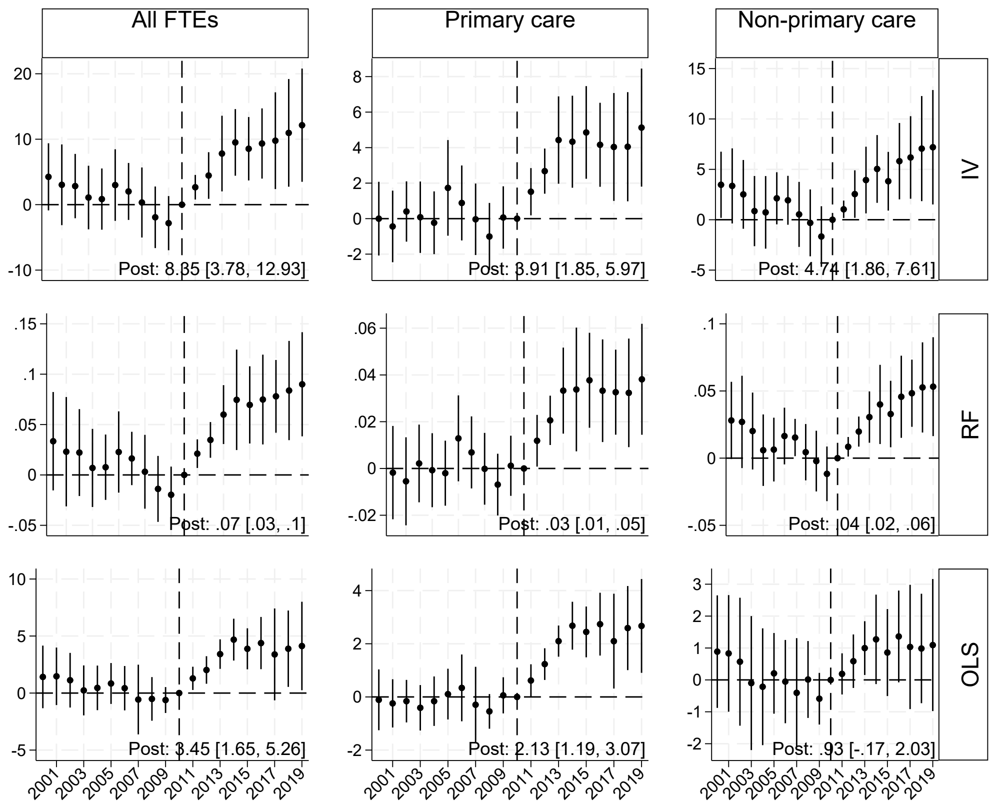

Source: Authors' analysis of CMS Healthcare Cost Report Information Systems hospital cost reports and Area Health Resource Files. *Notes*: Black dots in the first, second, and third rows correspond to point estimates for the λ_t coefficients in equations (3), (4), and (1), respectively. Solid vertical lines correspond to the 95 percent confidence interval from standard errors clustered at the state level. Regressions weighted by hospital beds. Average of post-period coefficients provided in the bottom right of each panel.

As a falsification exercise, we estimate the hospital-level models using as outcome variables various non-physician measures of hospital employment. In particular, we follow Prager and Schmitt (2021) and aggregate measures of total employee hours worked for narrow employment categories from the cost report data into hours for unskilled, skilled, and nursing and pharmaceutical employees. We then compute FTE counts for these three categories assuming a 40-hour workweek and use these to compute FTEs per 100 beds.[15]

[15] The formula used in Prager and Schmitt (2021) to convert employee category i hours for hospital h in

Results from estimating our hospital-level models using the binary indicator for any cap increase as the endogenous treatment variable and the non-resident FTE counts per 100 beds as our outcome variable are presented in Figure A5. The results show no impact of the cap increase on any of these non-physician measures of employment .

5.2 Residency program size

Figure 8 plots the event-study coefficients for the estimation of reduced form equation 4 using program-level data on Match outcomes. We are limited to estimating the reduced form model as we are unable to identify which programs are affiliated with treated teaching hospitals. The figures in the left column provide results for primary care programs, while those in the right column are for non-primary care programs. The outcome variable for results provided in the top row is the program's quota in the Match, while the outcome for results provided in the bottom row is the number of filled slots. The results indicate that in the years following Section 5503's implementation, a primary care program being located in a state one spot lower in the resident-to-population ratio distribution is associated with an increase in the program's annual quota of 0.12 and in its number of matched slots of 0.15. This implies that moving from the top to the bottom quartile of the resident-to-population ratio distribution is associated with an increase in a program's quota and matches of 1.56 (18%) and 1.95 (23%), respectively. We find no statistically significant effects for non-primary care programs. Together, the program-level results presented here and the hospital-level results presented in Section 5.2 indicate that Section 5503 was successful at increasing the number of residents recruited into primary care specialties as well as time spent at hospitals in high-need areas.

year t to FTEs is $\text{FTEs}_{iht} = \frac{365}{\text{CostReportDays}_{ht}} \times \frac{\text{TotalHours}_{iht}}{52 \times 40}$ where $\text{CostReportDays}_{ht}$ is the number of days covered by the cost report. Unskilled employee hours include hours for the Maintenance & Repairs, Operation of Plant, Laundry & Linen Service, Housekeeping, Dietary, Cafeteria, Central Services & Supply, and Medical Records & Medical Records Library line items. Skilled employee hours include hours for the Employee Benefits Department, Administrative & General, Maintenance of Personnel, and Social Service line items. Nursing and pharmaceutical employee hours include hours for the Nursing Administration and Pharmacy line items.

Figure 8: Estimation results for the effect of Section 5503 on program Match outcomes

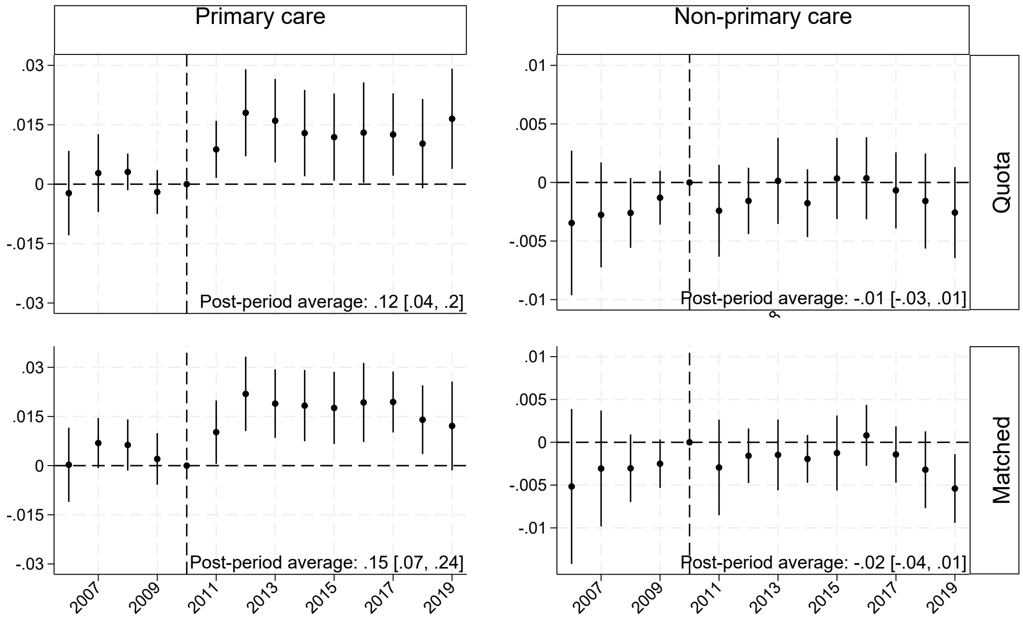

Source: Authors' analysis of National Residency Matching Program reports. *Notes*: Black dots correspond to point estimates for the λ_t coefficients in equation (4). Solid vertical lines correspond to the 95 percent confidence interval from standard errors clustered at the state level. Average of post-period coefficients provided in the bottom right of each panel.

5.3 Primary care physician supply

Figure 9 presents the results of estimating our county-level models. The outcome variable for these models is the natural logarithm of PCP counts for all PCPs, allopathic PCPs (aka MDs), and osteopathic PCPs (aka DOs) in the left, middle, and right columns, respectively. All estimates have been multiplied and scaled by a factor of 100. The estimates shown are therefore of the percent change in PCP supply affected by Section 5503 in the given year. The IV estimates show that Section 5503 increased PCP supply by 5.2 percent in treated counties, all of which is attributable to a 5.8 percent increase in primary care MDs. Results for DOs are not statistically significant and show evidence of pre-trends across specifications. OLS estimates also show statistically significant increases in primary care physician and MD supply and statistically insignificant effects for DO supply. The OLS estimates are smaller in magnitude than those for the IV model, indicating that treatment results in an average increase in PCP supply following Section 5503 of 1.9 percent for all PCPs and 2.3 percent for primary care MDs.

Figure 9: Estimation results for the percent effect of Section 5503 on county primary care physician counts

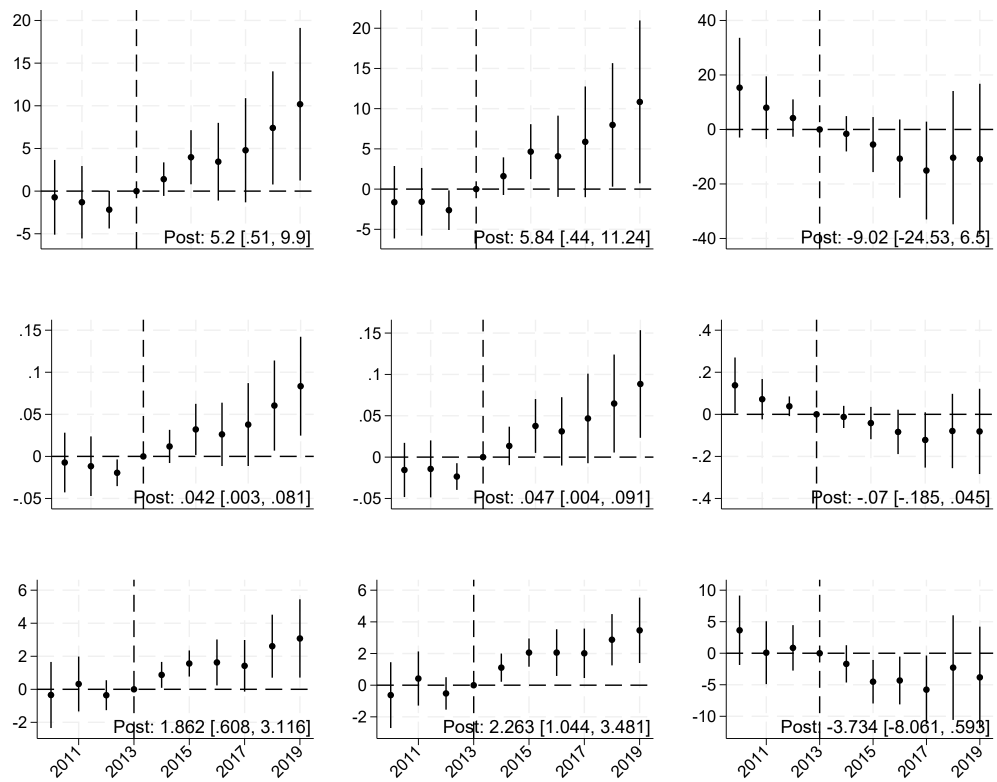

Source: Authors' analysis of CMS Healthcare Cost Report Information Systems hospital cost reports and Area Health Resource Files. *Notes:* Black dots in the first, second, and third rows correspond to point estimates for the λ_t coefficients in equations (3), (4), and (1), respectively. All estimates have been scaled by a factor of 100. Solid vertical lines correspond to the 95 percent confidence interval from standard errors clustered at the state level. Regressions weighted by county population. Average of post-period coefficients provided in the bottom right of each panel.

Figures A6 and A7 provide IV and OLS estimates for specifications in which the endogenous treatment variable is the cumulative number of slots awarded and the natural logarithm of the increase in GME subsidies, respectively, rather than the binary indicator for containing at least one treated teaching hospital used in the specifications in Figure 9. These results indicate that an increase in a county's number of subsidized residency slots of

10 yields an increase in PCP supply in the medium run of 1.3 percent and that a 10 percent increase in a county's cumulative GME funding increases PCP supply by 3.4 percent. Results using hospital referral regions rather than counties as a unit of observation yield qualitatively similar results.

6 Perfect conversion counterfactual

The estimates above show that Section 5503 yielded increases in the number of medical school graduates specializing in primary care and the time spent by primary care trainees at underserved hospitals. We can assess whether having trainees train in underserved areas is an effective way of influencing them to practice in those areas in the long run by comparing our estimated effects of Section 5503 on PCP supply to the effect we would see under perfect conversion. By "perfect conversion," we mean that for every resident recruited to a teaching hospital as a result of Section 5503, there is a corresponding increase in attending PCP supply of one physician in the county of that teaching hospital from the time of that resident's graduation through the end of our sample period.

Table 4 demonstrates how the annual hospital-level effects on primary care resident utilization from Section 5.1 are used to compute the cumulative effect on PCP supply per 100 beds under the assumption of perfect conversion. Column (1) provides the post-period point estimates plotted in the top middle panel of Figure 7. In column (2), these are lagged by two years to take into account the fact that residents recruited in year t take three years to complete primary care residency and will enter the pool of attending PCPs in year $t+2$. In column (3), we compute the running total of recruited PCPs to get the cumulative effect on PCP supply per 100 beds. We then multiply the sums in (3) by each treated county's number of treated beds and divide this by the corresponding county's baseline number of PCPs to get the percent increase in PCP supply under perfect conversion. We take the population-weighted average of these percent increases across counties within a year and present these averages in column (4). The average of these percent changes over the period 2013 to 2019 is 16.6. This is substantially larger than the average of our estimated annual effects from Section 5.3, which suggests that the conversion of residents recruited under Section 5503 was not perfect.

We solve for the conversion rate $\phi \in [0, 1)$ such that scaling the point estimates in column

Table 4: Computation of cumulative effect of Section 5503 on PCP supply per 100 beds under perfect conversion using estimates from Figure 7

Year	(1) Point estimate	(2) Point estimate of $t-2$	(3) Cumulative sum of (2)	(4) % change in PCP supply
2011	1.517	0	0	0
2012	2.676	0	0	0
2013	4.427	1.517	1.517	1.0
2014	4.333	2.676	4.193	5.2
2015	4.856	4.427	8.620	10.8
2016	4.160	4.333	12.953	16.2
2017	4.038	4.856	17.809	22.2
2018	4.051	4.160	21.969	27.5
2019	5.127	4.038	26.007	32.5

Source: Authors' analysis of CMS Healthcare Cost Report Information Systems hospital cost reports and Area Health Resource Files. *Notes*: Point estimates in column (1) correspond to the event study coefficients from the estimation of the hospital-level IV model summarized by equations (2) and (3) where the outcome variable is primary care FTE residents per 100 beds. These estimates are are plotted in the top middle panel of Figure 7. Values in column (4) computed as the population-weighted average of the product of values in column (3) by county-level number of treated hospital beds, divided by baseline PCP supply.

(1) by ϕ results in an average effect on PCP supply across years equal 5.2, our IV estimate from Section 5.3. Doing so yields $16.6\phi = 5.2 \implies \phi = 0.31$. ϕ can be interpreted as the share of residents recruited to underserved teaching hospitals under Section 5503 that are converted into an attending PCP practicing in the county of that teaching hospital in the medium run. This conversion may occur through retention - i.e. it may be the recruited resident themselves who decides to stay and practice in the area of their residency - or through attending PCPs from elsewhere as complements to recruited residents. This conversion rate of 0.31 is rather low, suggesting the causal effect of place of training on the location choices of physicians is modest.

7 Conclusion

In this study, we evaluate targeted subsidies for medical training as a means of addressing physician shortages. We estimate the effect of a change to the formula by which teaching hospitals in rural counties and states with low PCP supply were reimbursed for training residents. We find that this revision successfully increased treated hospitals' demand for residents, resulting in an increase in teaching intensity at teaching hospitals of per 100 beds. This increase in residents trained in high-need areas, in turn, resulted in an increase in attending PCPs practicing in those areas of 5.2 percent. Our results additionally imply that 31 percent of residents recruited as a result of Section 5503 were converted into attending PCPs practicing in the same county.

Other interventions such as Loan Forgiveness Programs (LFPs) (Falcettoni, 2017; Kulka and McWeeny, 2019) and the Conrad 30 Program (Braga et al., 2023) have been successful in attracting physicians to rural and underserved areas through financial or immigration incentives. Section 5503 differs from these interventions in that it indirectly addresses physician shortages by targeting hospitals' demand for trainees. A natural question is how the cost per additional PCP recruited to an underserved hospital under Section 5503 compares to that for LFPs. Kulka and McWeeny (2019) estimate that LFPs increased the number of PCPs in treated counties by 1.5 in the first year of the policy, with a median forgiven loan of \$100k. Section 5503 increased the number of PCPs per treated county by 5.2 percent. Given a baseline average of approximately 400 PCPs per treated county, this corresponds to an increase of 20 PCPs per treated county, each at a cost of approximately \$100k. Not only is Section 5503 more effective than LFPs at increasing PCP supply in underserved areas for a similar cost per added PCP, but it has the added benefit of doing so by increasing the overall supply of PCPs rather than reallocating PCPs from urban to rural areas, as is shown by our program-level results. Medical students trained in rural areas might also be expected to provide higher quality care than PCPs lured there by the promise of loan forgiveness. However, our results do not allow us to speak to this quality issue or, more generally, to the characteristics of the marginal resident recruited as a result of these programs. Quantifying these important effects will require panel micro-data on resident characteristics and employment decisions, which we leave to future work.

References

AAMC (2021a): "The Complexities of Physician Supply and Demand: Projections From 2019 to 2034," https://www.aamc.org/media/54681/download.

——— (2021b): "Report on Residents," https://www.aamc.org/data-reports/students-residents/data/report-residents/2021/executive-summary.

Angrist, J., D. Autor, and A. Pallais (2022): "Marginal Effects of Merit Aid for Low-Income Students," *The Quarterly Journal of Economics*, 137, 1039–1090.

Angrist, J., D. Lang, and P. Oreopoulos (2009): "Incentives and Services for College Achievement: Evidence from a Randomized Trial," *American Economic Journal: Applied Economics*, 1, 136–163.

Arcidiacono, P. and S. Nicholson (2005): "Peer effects in medical school," *Journal of Public Economics*, 89, 327–350.

Basu, S., S. A. Berkowitz, R. L. Phillips, A. Bitton, B. E. Landon, and R. S. Phillips (2019): "Association of Primary Care Physician Supply With Population Mortality in the United States, 2005-2015," *JAMA Internal Medicine*, 179, 506–514.

Basu, S., R. S. Phillips, S. A. Berkowitz, B. E. Landon, A. Bitton, and R. L. Phillips (2021): "Estimated Effect on Life Expectancy of Alleviating Primary Care Shortages in the United States," *Annals of Internal Medicine*, 174, 920–926.

Bazzoli, G. J. (1985): "Does educational indebtedness affect physician specialty choice?" *Journal of Health Economics*, 4, 1–19.

Braga, B., G. Khanna, and S. Turner (2023): "Migration Policy and the Supply of Foreign Physicians: Evidence from the Conrad 30 Waiver Program," NBER Working Paper 32005.

CMS (2018): "Federal Support for Graduate Medical Education: An Overview," https://healthcareworkforce.georgia.gov/main-publications-reports/data-publications/gbpw-fact-sheets/graduate-medical-education-fact-sheets.

Cohen, J. W. and W. D. Spector (1996): "The effect of Medicaid reimbursement on quality of care in nursing homes," *Journal of Health Economics*, 15, 23–48.

Currie, J. and J. Gruber (1996): "Health Insurance Eligibility, Utilization of Medical Care, and Child Health," *The Quarterly Journal of Economics*, 111, 431–466.

Dahl, G. B. and L. Lochner (2012): "The Impact of Family Income on Child Achievement: Evidence from the Earned Income Tax Credit," *American Economic Review*, 102, 1927–56.

Deming, D. J., J. S. Hastings, T. J. Kane, and D. O. Staiger (2014): "School Choice, School Quality, and Postsecondary Attainment," *American Economic Review*, 104, 991–1013.

Diamond, R. (2016): "The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980-2000," *American Economic Review*, 106, 479–524.

Fadlon, I., F. P. Lyngse, and T. H. Nielsen (2020): "Early Career Setbacks and Women's Career-Family Trade-Off," Tech. Rep. w28245, National Bureau of Economic Research, Cambridge, MA.

Falcettoni, E. (2017): "The Determinants of Physicians' Location Choice: Understanding the Rural Shortage," Working paper.

Feng, Z., D. C. Grabowski, O. Intrator, J. Zinn, and V. Mor (2008): "Medicaid Payment Rates, Case-Mix Reimbursement, and Nursing Home Staffing: 1996-2004," *Medical Care*, 46, 33–40.

Fordyce, M. A., M. P. Doescher, F. M. Chen, and L. G. Hart (2012): "Osteopathic physicians and international medical graduates in the rural primary care physician workforce," *Family Medicine*, 44, 396–403.

Foster, A. D. and Y. S. Lee (2015): "Staffing subsidies and the quality of care in nursing homes," *Journal of Health Economics*, 41, 133–147.

GAO (2017): "Locations and Types of Graduate Training Were Largely Unchanged, and Federal Efforts May Not Be Sufficient to Meet Needs," `https://www.gao.gov/products/gao-17-411`.

——— (2021): "Caps on Medicare Funded Graduate Medical Education at Teaching Hospitals," `https://www.gao.gov/assets/gao-21-391.pdf`.

Georgia Student Finance Commission (2023): "Teacher Education Assistance for College and Higher Education (TEACH) Grants," https://www.gafutures.org/federal-aid-scholarships/federal-student-grants/teach/.

Goodman-Bacon, A. (2021): "The Long-Run Effects of Childhood Insurance Coverage: Medicaid Implementation, Adult Health, and Labor Market Outcomes," *American Economic Review*, 111, 2550–2593.

GPBW (2007): "Fact Sheet on Georgia's Teaching Hospitals: Training Physicians and Providing Healthcare to Georgians," file:///Users/cici/Downloads/Fact%20Sheet%20-%20GA's%20Residency%20Programs%202007.pdf.

Greenwood, M. J. (1997): "Internal Migration in Developed Countries," in *Handbook of Population and Family Economics*, ed. by M. R. Rosenzweig and O. Stark, New York: Elsevier, 647–720.

He, D., P. McHenry, and J. M. Mellor (2020): "The Effects of Medicare Payment Changes on Nursing Home Staffing," *American Journal of Health Economics*, 6, 411–443.

HHS (2022): "New Reports Show Record 35 Million People Enrolled in Coverage Related to the Affordable Care Act, with Historic 21 Million People Enrolled in Medicaid Expansion Coverage," https://www.hhs.gov/about/news/2022/04/29/new-reports-show-record-35-million-people-enrolled-in-coverage-related-to-the-affordable-care-act.html.

Holzer, H. J., R. N. Block, M. Cheatham, and J. H. Knott (1993): "Are Training Subsidies for Firms Effective? The Michigan Experience," *Industrial and Labor Relations Review*, 46, 625–636.

HRSA (2022): "Primary Care Workforce: Projections, 2020-2035," https://bhw.hrsa.gov/sites/default/files/bureau-health-workforce/Primary-Care-Projections-Factsheet.pdf.

Kaestner, R. and J. Guardado (2008): "Medicare reimbursement, nurse staffing, and patient outcomes," *Journal of Health Economics*, 27, 339–361.

KFF (2023): "Primary Care Health Professional Shortage Areas (HPSAs)," https://www.kff.org/other/state-indicator/primary-care-health-professional-shortage-areas-hpsas.

Konetzka, R. T., D. Yi, E. C. Norton, and K. E. Kilpatrick (2004): "Effects of Medicare Payment Changes on Nursing Home Staffing and Deficiencies," *Health Services Research*, 39, 463–488.

Ku, B. S. and B. G. Druss (2020): "Associations Between Primary Care Provider Shortage Areas and County-Level COVID-19 Infection and Mortality Rates in the USA," *Journal of General Internal Medicine*, 35, 3404–3405.

Kulka, A. and D. McWeeny (2019): "Rural Physician Shortages and Policy Intervention," Working Paper.

Lin, H. (2014): "Revisiting the relationship between nurse staffing and quality of care in nursing homes: an instrumental variables approach," *Journal of Health Economics*, 37, 13–24.

Macinko, J., B. Starfield, and L. Shi (2007): "Quantifying the health benefits of primary care physician supply in the United States," *International Journal of Health Services: Planning, Administration, Evaluation*, 37, 111–126.

Markowitz, S. and K. Adams (2022): "The Effects of State Scope of Practice Laws on the Labor Supply of Advanced Practice Registered Nurses," *American Journal of Health Economics*, 8, 65–98.

McKay, N. L. (1990): "The economic determinants of specialty choice by medical residents," *Journal of Health Economics*, 9, 335–357.

McNamara, C. and T. Hussain (2023): "Graduate Medical Education Funding and Resident Utilization," Working paper.

MedPAC (2001): "Training requirements and initial residency period for selected residency programs," https://www.medpac.gov/wp-content/uploads/import_data/scrape_files/docs/default-source/reports/Mar01AppC.pdf.

Milligan, K. and M. Stabile (2011): "Do Child Tax Benefits Affect the Well-Being of Children? Evidence from Canadian Child Benefit Expansions," *American Economic Journal: Economic Policy*, 3, 175–205.

Nicholson, S. (2002): "Physician Specialty Choice under Uncertainty," *Journal of Labor Economics*, 20, 816–847.

Nicholson, S. and D. Song (2001): "The incentive effects of the Medicare indirect medical education policy," *Journal of Health Economics*, 20, 909–933.

NRMP (2010): "Results and Data: 2010 Main Residency Match," https://www.nrmp.org/wp-content/uploads/2021/07/resultsanddata2010.pdf.

——— (2022): "Results and Data: 2022 Main Residency Match," https://www.nrmp.org/wp-content/uploads/2022/11/2022-Main-Match-Results-and-Data-Final-Revised.pdf.

——— (2024): "Match Data," https://www.nrmp.org/.

Panzirer, W. (2021): "Survey Confirms Effectiveness Of Telehealth In Rural America And Beyond," https://www.healthaffairs.org/content/forefront/survey-confirms-effectiveness-telehealth-rural-america-and-beyond.

Prager, E. and M. Schmitt (2021): "Employer Consolidation and Wages: Evidence from Hospitals," *American Economic Review*, 111, 397–427.

Sanderson, E. and F. Windmeijer (2016): "A weak instrument F-test in linear IV models with multiple endogenous variables," *Journal of Econometrics*, 190, 212–221.

Shi, L., B. Starfield, B. Kennedy, and I. Kawachi (1999): "Income inequality, primary care, and health indicators," *The Journal of Family Practice*, 48, 275–284.

SSA (2007): "Nation's First Baby Boomer Files for Social Security Retirement Benefits – Online!" https://www.ssa.gov/pressoffice/pr/babyboomerfiles-pr.htm.

Starfield, B. (1991): "Primary Care and Health: A Cross-National Comparison," *JAMA*, 266, 2268–2271.

Thornton, J. (2000): "Physician choice of medical specialty: do economic incentives matter?" *Applied Economics*, 32, 1419–1428.

Thornton, J. and F. Esposto (2003): "How important are economic factors in choice of medical specialty?" *Health Economics*, 12, 67–73.

Wasserman, M. (2023): "Hours Constraints, Occupational Choice, and Gender: Evidence from Medical Residents," *The Review of Economic Studies*, 90, 1535–1568.

White, C. (2005): "Medicare's Prospective Payment System for Skilled Nursing Facilities: Effects on Staffing and Quality of Care," *Inquiry*, 42, 351–366.

Appendix A Figures

Figure A1: Trend in resident hourly real wage for hospitals that received cap increases and control hospitals

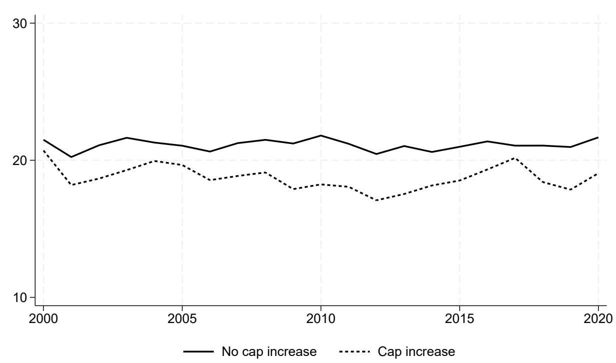

Source: Authors' analysis of CMS Healthcare Cost Report Information Systems hospital cost reports and price index data from the Bureau of Labor Statistics. *Notes:* Solid line shows the trend in the average hourly real wage for residents of teaching hospitals that did not receive a cap increase under Section 5503, while the dashed line shows the same trend for residents of teaching hospitals that did receive a cap increase under Section 5503.

Figure A2: Trend in distribution of hourly real wages

(a) All hospital employees (b) Residents

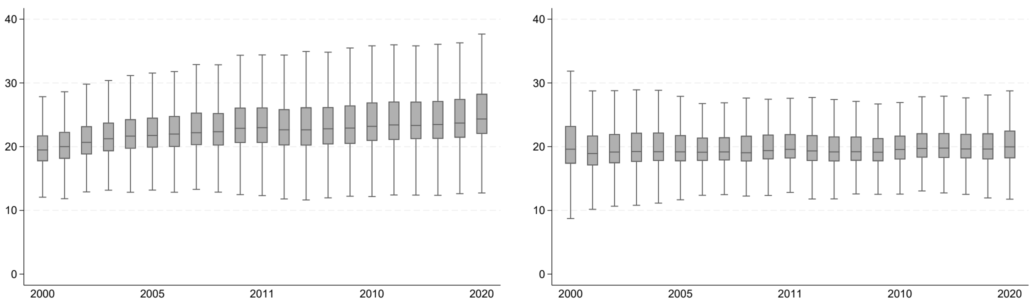

Source: Authors' analysis of CMS Healthcare Cost Report Information Systems hospital cost reports and price index data from the Bureau of Labor Statistics. *Notes:* Figures show the trend in the distribution of hourly real wages for residents at teaching hospitals and all other hospital employees at teaching hospitals.

Figure A3: Estimation results for the effect of a one slot increase in resident cap from Section 5503 on hospital residents per 100 beds

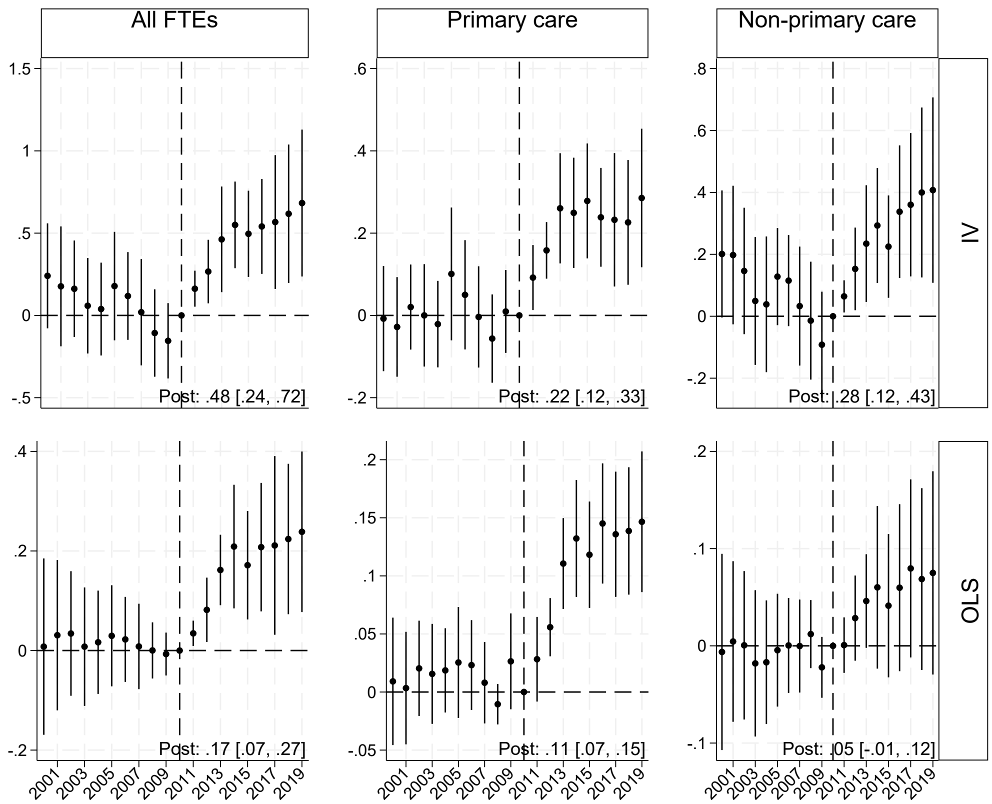

Source: Authors' analysis of CMS Healthcare Cost Report Information Systems hospital cost reports and Area Health Resource Files. *Notes*: Black dots in the first, second, and third rows correspond to point estimates for the λ_t coefficients in equations (3), (4), and (1), respectively. The endogenous treatment variable in these specifications is the number of residency slots awarded to the hospital. Solid vertical lines correspond to the 95 percent confidence interval from standard errors clustered at the state level. Regressions weighted by hospital beds. Average of post-period coefficients provided in the bottom right of each panel.

Figure A4: Estimation results for the effect of a one percent increase in residency subsidies from Section 5503 on hospital residents per 100 beds

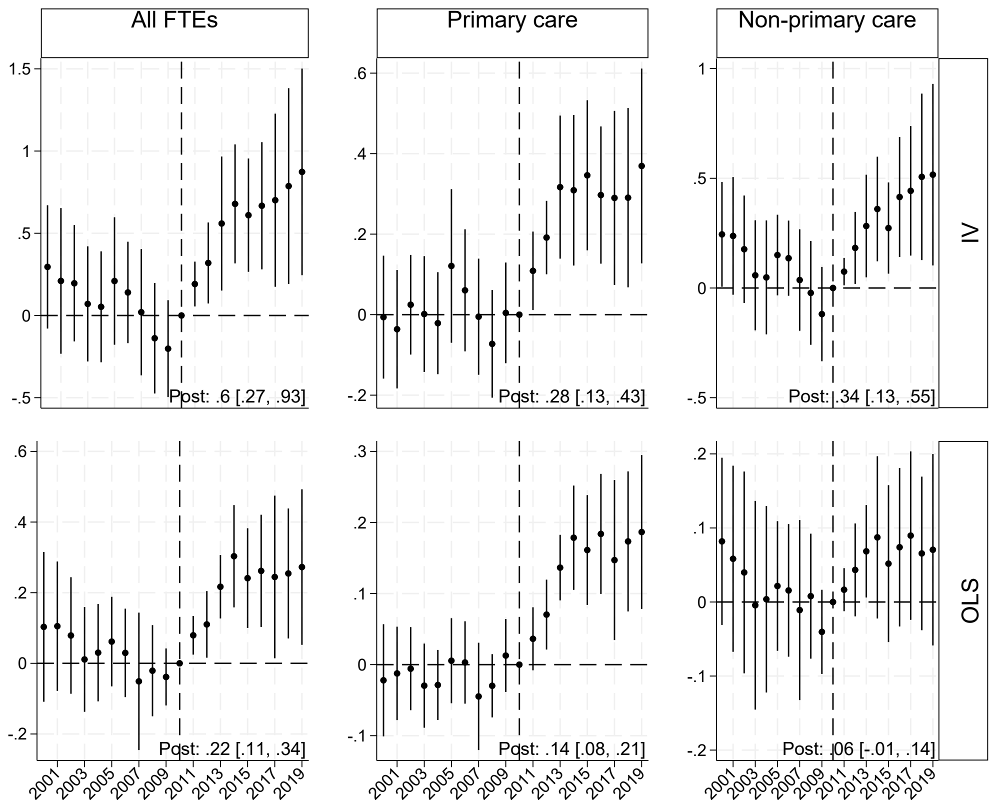

Source: Authors' analysis of CMS Healthcare Cost Report Information Systems hospital cost reports and Area Health Resource Files. *Notes*: Black dots in the first, second, and third rows correspond to point estimates for the λ_t coefficients in equations (3), (4), and (1), respectively. The endogenous treatment variable in these specifications is the natural logarithm of the potential GME subsidy to be earned from the residency slots awarded to the hospital. Solid vertical lines correspond to the 95 percent confidence interval from standard errors clustered at the state level. Regressions weighted by hospital beds. Average of post-period coefficients provided in the bottom right of each panel.

Figure A5: Estimation results for the effect of Section 5503 on other hospital employment per 100 beds

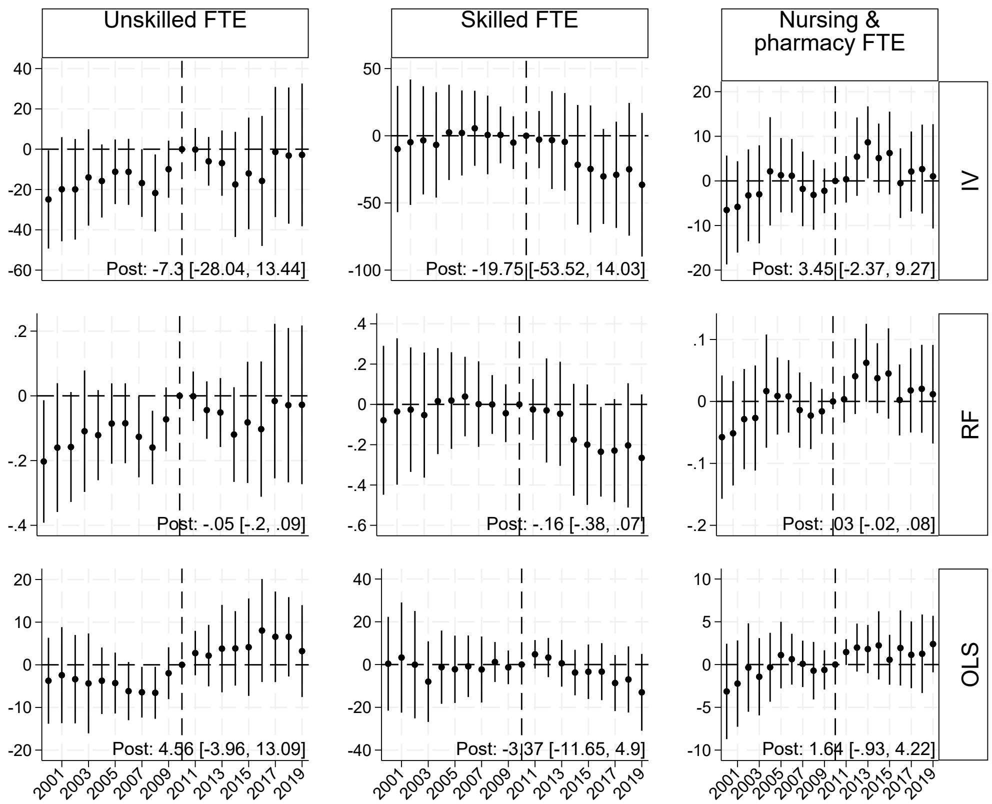

Source: Authors' analysis of CMS Healthcare Cost Report Information Systems hospital cost reports and Area Health Resource Files. *Notes*: Black dots in the first, second, and third rows correspond to point estimates for the λ_t coefficients in equations (3), (4), and (1), respectively. Solid vertical lines correspond to the 95 percent confidence interval from standard errors clustered at the state level. Average of post-period coefficients provided in the bottom right of each panel.

Figure A6: Estimation results for the percent effect of a one slot increase in the cumulative resident cap on county primary care physician counts

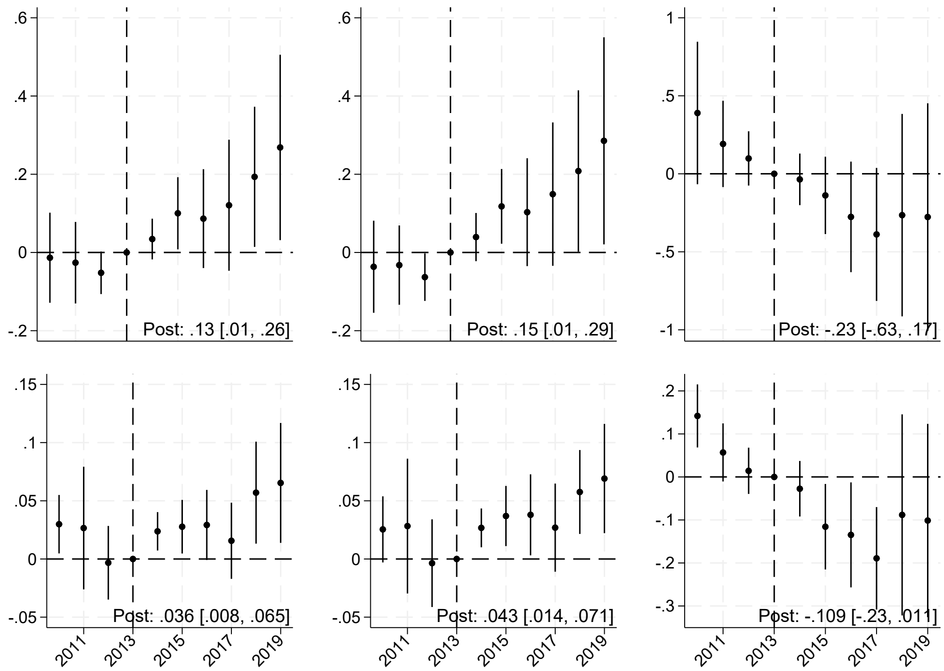

Source: Authors' analysis of CMS Healthcare Cost Report Information Systems hospital cost reports and Area Health Resource Files. *Notes:* Black dots in the first, second, and third rows correspond to point estimates for the λ_t coefficients in equations (3), (4), and (1), respectively. The endogenous treatment variable in these specifications is the sum of residency slots awarded to all hospitals in the county. All estimates have been scaled by a factor of 100. Solid vertical lines correspond to the 95 percent confidence interval from standard errors clustered at the state level. Regressions weighted by county population. Average of post-period coefficients provided in the bottom right of each panel.

Figure A7: Estimation results for the percent effect of a one percent increase in residency subsidies from Section 5503 on county primary care physician counts

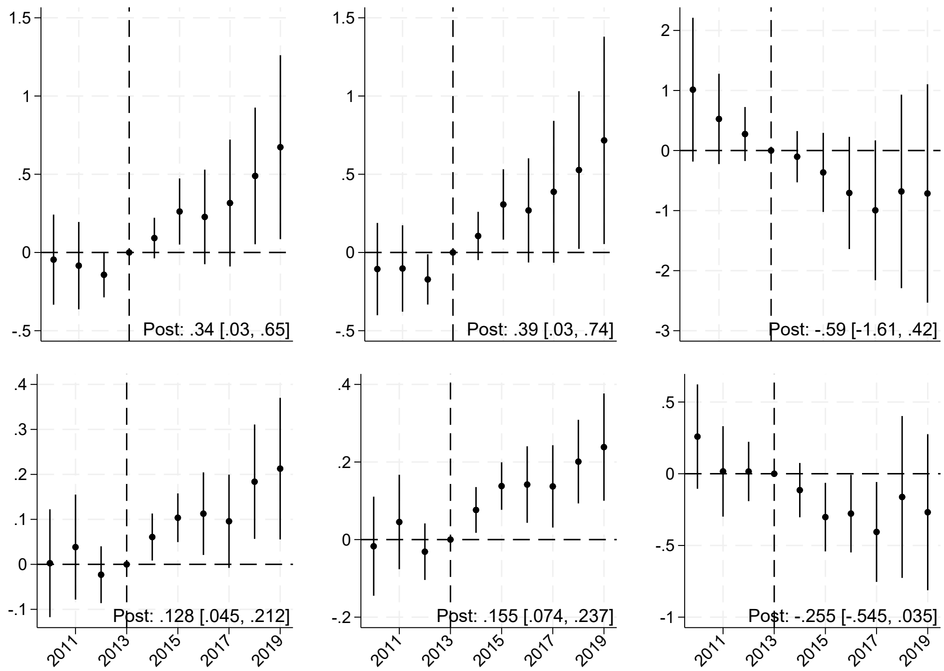

Source: Authors' analysis of CMS Healthcare Cost Report Information Systems hospital cost reports and Area Health Resource Files. *Notes:* Black dots in the first, second, and third rows correspond to point estimates for the λ_t coefficients in equations (3), (4), and (1), respectively. The endogenous treatment variable in these specifications is the natural logarithm of the sum of potential GME subsidy to be earned from the residency slots awarded to all hospitals in the county. All estimates have been scaled by a factor of 100. Solid vertical lines correspond to the 95 percent confidence interval from standard errors clustered at the state level. Regressions weighted by county population. Average of post-period coefficients provided in the bottom right of each panel.

Appendix B Tables

Table B1: Summary of first stage estimates

Panel A: Hospital-level models

	Coefficient	Standard error	SW F-statistic
$\mathbb{1}\{\text{Increase}\}_h \times$			
$\mathbb{1}_{2000}$	0.00831	(0.00190)	316.7
$\mathbb{1}_{2001}$	0.00886	(0.00231)	118.8
$\mathbb{1}_{2002}$	0.00805	(0.00223)	176.8
$\mathbb{1}_{2003}$	0.00830	(0.00215)	142.0
$\mathbb{1}_{2004}$	0.00810	(0.00201)	117.4
$\mathbb{1}_{2005}$	0.00768	(0.00209)	193.7
$\mathbb{1}_{2006}$	0.00645	(0.00204)	233.3
$\mathbb{1}_{2007}$	0.00641	(0.00196)	151.9
$\mathbb{1}_{2008}$	0.00620	(0.00174)	264.7
$\mathbb{1}_{2009}$	0.00652	(0.00217)	168.4
$\mathbb{1}_{2011}$	0.00671	(0.00225)	209.6
$\mathbb{1}_{2012}$	0.00640	(0.00248)	84.2
$\mathbb{1}_{2013}$	0.00671	(0.00230)	381.1
$\mathbb{1}_{2014}$	0.00707	(0.00217)	475.6
$\mathbb{1}_{2015}$	0.00786	(0.00202)	241.5
$\mathbb{1}_{2016}$	0.00813	(0.00216)	178.7
$\mathbb{1}_{2017}$	0.00819	(0.00216)	252.6
$\mathbb{1}_{2018}$	0.00811	(0.00213)	176.9
$\mathbb{1}_{2019}$	0.00769	(0.00209)	159.0
Cragg-Donaldson F-statistic	76.48		
Kleibergen-Paap rk LM-statistic	3.951		

Panel B: County-level models

	Coefficient	Standard error	SW F-statistic
$\mathbb{1}\{\text{Increase}\}_c \times$			
$\mathbb{1}_{2000}$	0.00893	(0.00355)	66.3
$\mathbb{1}_{2011}$	0.00937	(0.00347)	31.5
$\mathbb{1}_{2012}$	0.00895	(0.00354)	42.3
$\mathbb{1}_{2014}$	0.00834	(0.00311)	28.9
$\mathbb{1}_{2015}$	0.00834	(0.00286)	38.8
$\mathbb{1}_{2016}$	0.00854	(0.00298)	34.6
$\mathbb{1}_{2017}$	0.00875	(0.00284)	29.3
$\mathbb{1}_{2018}$	0.00878	(0.00281)	19.3
$\mathbb{1}_{2019}$	0.00852	(0.00280)	23.9
Cragg-Donaldson F-statistic	138.7		
Kleibergen-Paap rk LM-statistic	4.097		

Source: Authors' analysis of CMS Healthcare Cost Report Information Systems hospital cost reports and Area Health Resource Files. *Notes*: We provide for each first-stage regression the sum of the estimated coefficients for the excluded instruments, the standard error of this combination of parameters, and the Sanderson and Windmeijer (SW) F-statistic. We also provide the Cragg-Donaldson F- and Kleibergen-Paap rk LM-statistics, which indicate that the instruments are strong and that the model is identified.