The Determinants of Physicians' Location Choice: Understanding the Rural Shortage*

Elena Falcettoni

Abstract

A long-standing challenge in the US health care system is the provision of medical services to rural areas. This paper develops a structural spatial equilibrium model with heterogeneous physicians and uses it to explore the impact of policies, namely loan forgiveness and salary incentives, on the geographical distribution of physicians. I collect micro data from physicians' directories on their medical school, residency, and first-job choices and use this new dataset to both allow for physicians' preference to remain close to their residency location and to implement an instrumental-variable approach to overcome endogeneity issues caused by the correlation between wages and unobserved amenities. I find that residents strongly prefer remaining close to their residency location. The combination of loan forgiveness and salary incentives has led to 1.2 percent more physicians choosing rural areas. Using the government spending currently allocated to loan forgiveness to further increase salary incentives would lead to 6 times more primary care physicians choosing rural areas and a higher average quality of rural physicians compared to the impact of current policies.

Keywords: rural physicians, healthcare, shortage, primary care, specialists, loan forgiveness.

*Board of Governors of the Federal Reserve System (email: elena.falcettoni@frb.gov). The views expressed herein are those of the Author's and do not represent the views of the Board of Governors of the Federal Reserve System or any person associated with the Federal Reserve System.
First Version: November 2017. The author is eternally grateful to Thomas Holmes and Amil Petrin for their guidance. The paper benefitted greatly from Robert Town's discussion during the Annual Health Econometrics Workshop 2020. My research was only made possible thanks to comments from Naoki Aizawa, Jenny Bourne, David Canning, Mariacristina De Nardi, Ben Handel, Karen Kopecky, Pinar Karaca Mandic, Corina Mommaerts, Hannah Neprash, James Schmitz Jr., Joel Waldfogel, and Motohiro Yogo, as well as participants of seminars at the University of Minnesota, Wake Forest University, the Federal Reserve Board of Governors, McGill University, Sciences Po, the University of New South Wales, the 2020 Southern Economics Association Annual meeting, the 2020 DC IO Day, the 2019 International Industrial Organization Conference, the 2019 and 2018 American Society of Health Economists Conferences, the 3rd Workshop on Mechanism Design for Social Good, the 2019 Urban Economics Association Conference, the 2018 Warwick PhD Economics Conference at Warwick University, the 2018 Midwest Economics Association Conference, and the 2017 MVEA conference. Funding from the University of Minnesota's Consortium on Law and Values in Health, Environment and the Life Sciences is gratefully acknowledged.

1 Introduction

A long-standing challenge in the U.S. health care system is the provision of medical services to medically-underserved areas, most of which rural, where 25 percent of the population live, but only 10 percent of physicians practice.[1] With roughly 60 million Americans living in rural areas, rural Americans make up for a major part of the affected population. Policymakers have tried to bring more physicians into rural areas, most notably with loan forgiveness programs, for which the United States roughly spends $350 million a year, and which have been further incentivized through their tax exclusion via the Affordable Care Act. As a further incentive, salary incentives are also usually offered to physicians who decide to practice in rural areas in addition to already-higher equilibrium wages. While the number of primary care physicians who practice rurally has increased, the medical association and policymakers continue to report an existing and worsening physician shortage, especially for some (rural) areas. Therefore, to understand how well these policies are faring (or failing,) it is first necessary to fully understand the different factors that affect physicians in their geographical choices to then be able to design policies that aim for a more even distribution of physicians in caring for the American population.

This paper develops a model of physicians' location choices and uses it to explore the impact of policy changes, such as loan forgiveness and salary incentives, on the geographical distribution of physicians. I focus on the choice of the first job following residency, therefore analyzing the location choice once the specialty is already picked. While studies have been done regarding the choice of residency, the so-called medical match (for instance, see Agarwal 2015,) this paper does not look at the specialty choice along with location. The decision to take the specialty as given allows me to ignore the issue of residency slots available as well as all the details of "the match." I build a structural spatial equilibrium model in which physicians are heterogeneous along their specialty, demographics, and quality ranking. Identification of the parameters of interest (income) is challenged by the possible correlation between unobserved characteristics of location and wages, as offered wages are higher where amenities are fewer. To overcome this issue, I collect micro-level data from physicians' directories on physicians' medical school, residency, and first-job choices. This wealth of information and structural methods of demand à la Berry, Levinsohn, and Pakes (1995) allow me to back up the unobserved characteristics and identify the parameters of interest exactly.

The differentiation between specialty groups is shown to be key in policy design. As shown in Falcettoni (2018), the mix of treatments that primary care physicians and specialists perform varies along the urbanity index. In particular, primary care physicians perform more specialty procedures (and therefore receive higher procedure revenue) in rural areas, due to the lower competition coming from specialists in those areas. Capturing this heterogeneity allows policies to be designed more efficiently by targeting the groups of physicians who would respond the most.

The information on physicians' quality ranking enables me to vary the job choice set according to each physician's quality level. This is important in the design of choice sets as the set of jobs available to each physician critically depends on the physician' quality. There is a vast literature addressing how misspecification of the choice sets leads to choice model misspecification. Gopinath (1995) provides a good overview of the theoretical and empirical issues on this topic. In my setup, a physician's quality is approximated by the specialty-specific quality ranking of that physician's medical school and residence. This is a fine approximation in this setting since I model the practice location for the physician's *first job* out of residency,

[1] See, among others, Aaron Carroll, "A Physician Shortage? Let's Take a Closer Look," *The New York Times*, November 7, 2016 as well as Gary Hart's interview, Ann Harrington, "Training More Country Physicians," *Fedgazette*, October 12, 2017.

so that educational and training attainment is the best proxy for that physician's skills in the eyes of every employer.

Several factors affect a physician's location choice. I allow physicians to respond to their net real income (therefore accounting for salary, procedure revenue, rent, malpractice insurance, and student loan repayment,) as well as to amenities and heterogeneous location preferences. One immediate trade-off related to physicians' incentives is the higher salary offered in rural areas to compensate individuals for the typical lack of amenities (see, for example, Lee 2010.) However, the urbanity of the area not only influences the amenities of the area but also the competition from physicians in surrounding areas, which in turn affects the procedures that physicians can carry out. Physicians' revenue stream is composed of two parts: a salary part that behaves as theory would predict by increasing salaries in less desired areas, and a procedure revenue part. This latter part depends only on the procedures carried out and is adjusted for the cost of living, meaning that the rate adjustments are greater for physicians in urban areas than for those in rural areas. Moreover, the number and type of services provided also vary along the urbanity index. Falcettoni (2018) shows that primary care physicians in rural areas are able to increase their income by carrying out more specialized procedures. Since the fee-for-service part of their income does not depend on their specialty but only on the procedures they carry out, and since such specialized procedures are increasingly-more remunerative than primary care procedures, this creates incentives for primary care physicians to work in rural areas. Therefore, these different components of income and how they are affected by the distribution of physicians must be accounted for. On the supply side, I allow physicians' income to respond to the employment of physicians of either type, bearing these facts in mind.

I also allow for a home bias toward the place where the physician completed his or her residency, based on data evidence. To be able to control for quality, I match the ranking of the medical school (based on the average score of MCATs, among other things) and the sub-specialty-specific ranking of the residency to proxy for physicians' quality ranking. As mentioned beforehand, quality not only is important as a demographic variable but also is key in the choice set definition.

Next, I analyze what factors affect their geographical distribution the most. Unsurprisingly, all physicians enjoy higher net incomes and higher amenities. I find that the two specialty groups respond to compensation differently, as specialists are more elastic to both net income and amenities. I find that top-50 residents respond more to both income and amenities, while foreign physicians are not systematically different from Americans. Persistence in location choice is key, as I find that primary care physicians are about 3.8 times more likely to pick a job within the same state of residency and about 3.4 times more likely to pick a job within the same hospital referral region as the residency. On the other hand, specialists are 2.8 times more likely to pick a job within the same state as residency and about 3.6 times more likely to pick a job within the same hospital referral region. I am able to reject that retention values can be the same between primary care physicians and specialists. I also find that top-50 residents in primary care are 0.4 times more likely to remain in the same state as residency, but they are 1.5 times less likely to remain in the same area as residency. On the other hand, I find that top-50 residents in specialty care are 0.3-0.4 times less likely to be retained within the same state and area of residency. Comparing these values to the corresponding home-state biases in the labor literature, I find a very interesting result. While all physicians are clearly high-skilled workers, primary care physicians display the same persistence in location choices as unskilled workers. Diamond (2015), for example, reports a semi-elasticity of retention for high-skilled individuals to remain within their state of birth equal to 2.6. That estimate is closer to the values I find for specialists, but much lower than the values that I find for primary care physicians. This shows that there are extremely important differences not only across occupation types, but also within occupations that might be ignored in current analyses.

Finally, I use the model to analyze the performance of current policies targeted at bringing physicians to rural areas. I find that the combination of loan forgiveness and salary incentives have led to 1.2 percent more physicians in rural areas. I find that 0.5 percent more primary care physicians and 1.3 percent more specialists have picked rural areas due to loan forgiveness alone. Monetary incentives in the form of bonus payments averaging $7,500 are responsible for a further 0.2 percent increase in primary care physicians and 0.1 percent increase in specialists. By retargeting the spending currently used for loan forgiveness to increase higher salary incentives for rural employment, I find that almost 6 times more primary care physicians would respond to that policy and pick rural areas because of the increased incentives compared to the effect generated by current policies. Since primary care physicians are the main physician category that currently provides medical care to rural areas, and since they do not need a particular infrastructure to do so, these results suggest that policymakers should retarget spending from loan forgiveness to salary incentives and that offering increased salary incentives as the main policy to attract primary care physicians to medically-underserved areas would be even more effective. The average quality of the physicians attracted under higher salary incentives is also better compared to the average quality of physicians attracted under the current policies. Finally, another possible policy intervention suggested by the high persistence in physicians' location choices is the use of these monetary incentives to create rural residencies. Since the residency choice is not directly modeled in this paper, this question is outside the scope of this paper but will be addressed in future work.

The paper proceeds as follows: Section 2 introduces a brief literature review, Section 3 presents a few definitions and descriptive facts, Section 4 describes the many data sources used in this paper, Section 5 examines the model, Section 6 discusses the estimation techniques, Sections 7 presents and discusses the results and their implications, Section 8 discusses the counterfactuals run, and Section 9 concludes.

2 Literature Review

This paper contributes to three strands of literature: it complements and extends the old microeconomics literature on physician location and geographical distribution, it provides more insight to the health economics literature on physicians' response to incentives, and it relates and extends the labor literature on location choice of skilled workers.

First, this paper contributes to the strand of literature on physician location. Cooper et al. (1975), Leonardson, Lapierre, and Hollingsworth (1985), Steele and Rimlinger (1965) are all papers that have provided evidence for an uneven distribution of physicians across the United States. These papers document this distribution and factors that impact physicians' choices of practice location through surveys and reduced-form analyses, but they do not provide a mechanism that explains these location choices. Previous discussion on the topic of location choice has mainly focused on the tradeoff between amenities and salary. Lee (2010) for example, provides evidence of higher salaries in rural areas relative to urban areas and provides a theory that the increased salary has to make up for the lack of amenities. There has also been a lot of attention regarding the shortage of physicians and the distribution of physicians' location, including Kirch, Henderson, and Dill (2012) and Cooper et al. (2002). This paper complements these analyses by identifying and studying the major factors that affect the choice of physicians' location and possible solution to the physicians' shortage.

Kulka and McWeeny (2018) also structurally analyzes physicians' location choices and evaluates policies the induce physicians to move to rural areas, but my analysis differs in several important respects. First, I differentiate across specialty groups. Second, I collect micro data on physicians' training and work history to estimate the value of retention from remaining within the same area as their residency and to define the choice set according to the physicians' quality. Finally, I employ a more detailed measure of compensation that includes net income that also depends on procedure revenue, rent, malpractice insurance, and student loan repayments. Importantly, the micro data that I collect on physician's education allows me to estimate student loans, which in turn allows me to directly estimate the impact of loan forgiveness policies and to compare this effect to the one of alternative policies on the distribution of physicians.

This paper also contributes to the strand of health economics literature discussing how physicians respond to financial incentives, basing part of the analysis on Falcettoni (2018). Falcettoni (2018) provides evidence for a supply-induced demand mechanism for more remunerative treatments. The paper finds that primary care physicians are able to take on more specialist services in less urban areas, where they gain higher market shares due to the lower number of specialists in close proximity. In particular, the increase in the weight of the primary care physicians' financial interests in the consumer utility ranges between 7-16 percent compared to physicians in large metropolitan areas, at the expense of specialists. More generally, Lee (2010) shows that higher rural salaries provide an incentive for physicians to trade off lower amenities for higher compensation. There has been an extensive literature on the response of physicians to financial incentives in a hospital setting (Acemoglu and Finkelstein 2008, Finkelstein 2007) , in managed care (Lori 2009), for specific procedures (Gruber & Owings 1994, Grant 2009, Shrank 2005, Jacobson 2006), and across geographical locations (Clemens and Gottlieb 2014). This paper complements this literature by including financial incentives in the analysis, without only focusing on wages, but also by analyzing how physicians trade off monetary incentives for non-monetary ones. Additionally, it also unveils the importance of including all financial incentives that can impact a physician's choice, from all sources of remuneration including procedure revenue to the monetary impact of policies such as loan forgiveness.

Since physicians are a large high-skill, high-income occupational group, this paper also complements the location choice strand of the labor literature across skill levels, including, but not limited to, Diamond (2015) and Colas (2018). Of course, physicians are all part of skilled labor. Nevertheless, this paper provides insight on within-occupation differences across types and shows that, at least for physicians, the differences between within-occupation types are just as important as those between the unskilled and the skilled. This paper also complements the structural labor literature on location choice and market structure, such as Dunne et al. (2013).

Methodologically, this paper bases itself mostly on Berry, Levinsohn, Pakes (1995, hereafter: BLP). While BLP has been one of the most predominant tools in the literature for demand estimation, this paper applies the tool to a location-choice setting. Thanks to the differentiation across locations and the presence of physician cohorts looking for a job at the same time nationally, I utilize this algorithm to identify what drives the choice of physicians' location, matching the share of physicians picking one location over all the physicians looking for a job in the same year. I include demographic characteristics of physicians and integrate over the empirical distribution of such characteristics to identify random coefficients. Two innovations are present: 1) I divide physicians into quality groups according to their education and training quality and allow top-ranked physicians to pick a job first, and 2) I build my instrumental variable in a way that resembles a Bartik share instrument, and therefore I do not use BLP instruments.

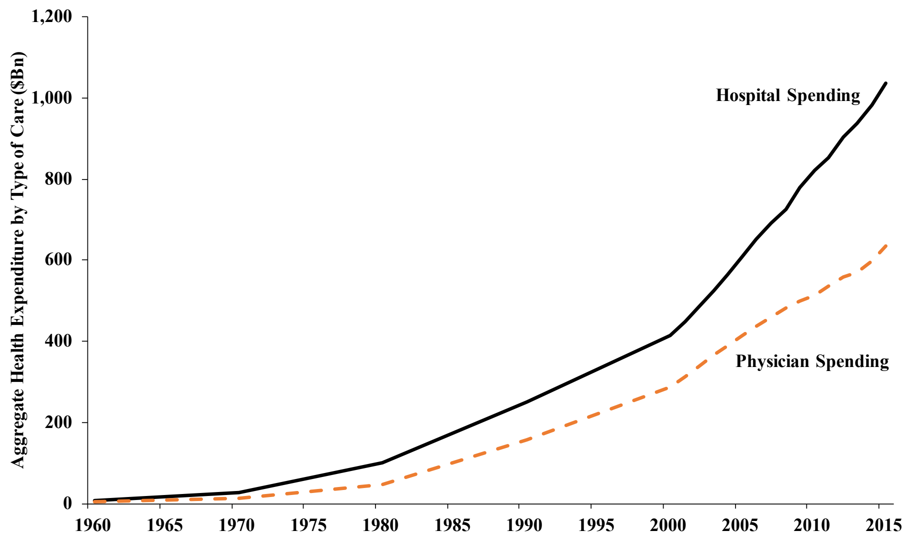

Figure 1: Aggregate Health Expenditure by Type of Care, $Bn

Notes: This figure shows the level of aggregate health spending differentiating between hospital care and physician-only care. Source: CMS.

3 Income Components

Income in my paper is calculated using five elements. First, I include both procedure revenue and salary to their total revenues. Second, I subtract from their total revenues three types of expenses: average housing cost in their area, which approximates for the local good cost, malpractice insurance payments, and student loan repayments.

Procedure revenue

Procedure revenue is the revenue stream that depends on the number and type of procedures carried out and not on the specialty of the physician carrying them out. Much of the health literature focuses on analyzing hospitals and physicians in a hospital setting. While physicians in my dataset can indeed have hospital affiliations, it is important to differentiate how I define their income. Since procedure revenue billed by hospitals is treated in a completely-different manner, I focus on aggregating procedure revenue from CMS Medicare Part B data, which excludes hospital bills. This is particularly important because common practice for hospitals is to file all procedure reimbursements for all physicians, and wages paid out to the physicians directly employed by the hospital will simply be adjusted for their corresponding procedure revenue, which would therefore be already accounted for in their salary component.

To be able to combine the effect of receiving procedure revenue and wages, I therefore single out procedure revenue billed for outpatient procedures, which are billed directly and paid out directly to the physicians. These actually make up for a high portion of aggregate health expenditures, as shown in Figure 1. Since they constitute a substantial portion of their income, physicians internalize them in their decision-making. I am able to identify whether a physician works at a hospital as well in two ways: 1) through the presence of an office at a hospital facility in CMS, and 2) through a declared hospital affiliation in the dataset I create by scraping data, as discussed below.

The current (since 1992) fee-for-service system is called the Resource-Based Relative Value Scale (RBRVS). The system was based on some initial rates and geographical adjustment factors, which would be reviewed on an annual basis by the RVS Update Committee (RUC). The RUC was meant to only have an advisory role, but its recommendations are accepted 97 percent of the time, making it *de facto* the fee-setting organization.

The Reader should bear in mind that the fee-for-service system is not new to 1992. The system before, the Usual, Customary, and Reasonable (UCR) system, was still based on a fee-for-service payment scheme; however, these rates were not standardized across physicians and tractability was not possible also due to lack of information on individual pricing. This is what prompted discussions at the beginning of 1990 to reform it. This standardization helps the estimation of the impact of this pricing on physicians' choices.

For each procedure j in a geographical area-year t, the procedure revenue is equal to:

$$ProcRev_{jt} = Constant_t * RVU_{jt} * GAF_t \tag{1}$$

The constant only depends on the year and is equal across specialties and procedures. The relative value units change according to the procedure as well as the year, and the geographic adjustment factors (GAFs) depend on both the area and the year.

The constant, called the Conversion Factor (CF), is a national adjustment factor, which is identical across specialties, areas, and procedures. The 2017 CF is equal to $35.8887. The GAFs are a proxy for cost of living, adjusting for differences in input costs across payment regions. This means that the GAFs make the per-procedure revenue higher in urban areas than in rural areas because of the higher cost of living in cities.

The RUC's recommendations across the years have been constantly widening the gap between the rates for procedures usually carried out by specialists and those regularly carried out by primary care physicians.

As shown in Falcettoni (2018), since these procedure rates do not depend on who carries out the procedure, but only on the procedure itself, and specialty procedures are more highly priced than typical primary care procedures, this payment system generates financial incentives for primary care physicians to substitute to more specialized, remunerative procedures whenever possible and this effect is stronger in rural areas. See Section A.1 in the Appendix for a longer discussion on these results.

Salaries

Salaries are obtained through the Bureau of Labor Statistics and do not include procedure revenue earned through procedures carried out in an office. BLS collects its data through a survey of employers that only report the salary paid. Any amounts billed independently by physicians is not included in the estimate. Self-employed physicians are not included in the wage estimates by BLS. I observe whether or not the physician is in a facility (hospital or clinic) or in an office setting and I am therefore able to differentiate between the two types of physicians in my analysis.

Housing costs

Housing costs are obtained from the Census American Community Survey, at the zip-code level. Average housing costs are calculated for individuals with incomes higher than $70,000. Notice that this income threshold is trivially crossed by all physicians considered in my analysis, since they have all completed their residency by the time they start their first job and income for a physician after residency is never below $70,000. I focus on owner costs and use the average between mortgage- and non-mortgage-holders.

Malpractice insurance

Malpractice insurance costs are estimated according to the malpractice reimbursement rates set by Medicare. In the future, I plan on adjusting these based on observed insurance rates.

Student loans

Student loan repayments are estimated in the following manner. First, I match each individual to the medical school he or she attended. Second, I match the medical school to the tuition cost for the four years of medical school, as available through their "Tuition and Rates" page online. Third, I follow the very common 10-year repayment plan most students would be on to calculate the average annual student loan repayment. According to the plan, the average interest rate is 6 percent, which is what I use in this paper. The interest starts accruing from year one, but payments are deferred until after residency, as it commonly happens.[2] For areas defined as health professional shortage areas, the cost of repayment is set to zero. Note that this assumes that residents who decide to move to these areas indeed remain there for the years necessary to have their loans forgiven. A separate project aims to dynamically estimate this year-by-year choice and its policy implications.

4 Data

4.1 Publicly-Available Data

The geographical unit of study is a hospital referral region (HRR), as defined by the Health Resources & Services Administration. Therefore, any location-level characteristics are estimated for these geographic areas through data at the county-, metro-, and zip-code level, allocated to HRRs by averaging or aggregating the values up to the HRR-level according to their geographical location. The goal of HRRs is to define areas that are self-contained markets for health care, so that the majority of patients living in that area go to physicians within that area.

[2] While residents face the choice to start repaying loans during residency, very few do.

The urban/rural classification follows the U.S. Census Bureau definition according to the 2010 Census criteria.

The primary source of data for this paper comes from the Centers for Medicare & Medicaid Services (CMS). The Physician and Other Supplier Public data provides information on services and procedures provided to Medicare beneficiaries by physicians. It contains information on utilization, actual Medicare procedure revenue, and submitted charges. Each line of the dataset is indexed by a National Provider Identifier (NPI), which identifies each physician in the dataset, by a Healthcare Common Procedure Coding System (HCPCS) code, which identifies every procedure carried out by each physician, and by the place of service, indicating whether the procedures were carried out in a facility setting or not. The data is based on information from CMS administrative claims data for Medicare beneficiaries enrolled in the fee-for-service program. The data are available for calendar years 2012 through 2016 and contain the universe of physicians taking part in Medicare Part B for the fee-for-service population. There are a little over 40 million observations in the dataset across over a million of physicians in the panel.

Despite the wealth of information on payment and utilization for Medicare Part B services, the dataset has a number of limitations. The data may not be representative of a physician's entire practice as it only includes information on Medicare fee-for-service beneficiaries. However, since Medicare influences the payment system of 80 percent of physicians (Clemens & Gottlieb 2017), these data allow for the analysis of physicians' behavior under this payment mechanism, which is then relevant for the greatest majority of physicians in the country. While private procedure rates can still differ substantially from Medicare rates, the paper will still produce qualitatively-correct estimates as long as there is no concentrated bias in the way rates differ across procedures or areas.

In addition, the data are not intended to indicate the quality of care provided and are not risk-adjusted to account for differences in underlying severity of disease of patient populations. To counter this, demographic data on patients' riskiness and incidence of diseases will be included in the estimation. As far as the illness and disease distribution is concerned, average beneficiary risk scores are provided on the "Medicare Physician and Other Supplier Aggregate Table" (i.e., one record per NPI) to provide information on the health status of the beneficiaries the providers serve for every year of interest together with the rate of incidence of a number of diseases and illnesses among the patients seen by each physician for every year. Therefore, this can account for the average health of the patients visited by each physician.

Despite these limitations, some positive characteristics should be highlighted. First of all, the fact that all beneficiaries are covered by Medicare eliminates the issues related to the status of insurance of the beneficiaries. In particular, it allows me to abstract from other endogenous characteristics related to the insurance status of beneficiaries when a full dataset (not Medicare only) is used. Moreover, it also allows me to ignore the network effects of different insurance policies as well as their different payment plans. In practice, therefore, this dataset provides a more homogenous universe of insured individuals who receive different treatments according to the condition that they have.

For this to work, it is important to recall that almost the universe of physicians participates in Medicare (and I will observe the physician as soon as they have one patient enrolled in Medicare Part B). To be precise, over 91 percent of physicians accept new Medicare patients and 96 percent of Medicare seniors have access to care through their physicians/clinic. Almost 92 percent of Medicare fee-for-service patients can get an appointment for routine care as needed.

The location characteristics included in the amenity index come from many different sources: County Business Patterns (2012-2016), Federal Bureau of Investigation crime reports (2012-2016), Environmental Pro-

Table 1: Physicians Work History Panel

	Primary Care	Specialists
Number of Residents	9,691	22,068
Years in the Panel	2012-2016	2012-2016
Locations Chosen Between 2012-6	305	305
percent with First Job in Big Metros	60 percent	70 percent
percent with First Job in Small Cities	30 percent	25 percent
percent with First Job in Rural Areas	10 percent	5 percent
percent that Completed Residency in Big Metros	58 percent	75 percent
percent that Completed Residency in Small Cities	31 percent	24 percent
percent that Completed Residency in Rural Areas	1 percent	1 percent

Notes: The Physician Work History Panel is a dataset that I created and that provides physician-level data on their training and on the work they currently carry out. The Panel is created through two main data sources: first, I use the Medicare Part B Utilization and Payment data; second, I scrape physician directories (mainly Doximity.com) to be able to determine their medical training (medical school, , residency). I then use the medical school information to infer the level of student debt they would be facing and I collect Bureau of Labor Statistics data on wages by occupation title to collect information on salaries. While this panel covers almost the universe of physicians, I focus on residents that finish residency and enter the medical job market in this paper. More details on the data collection and sample validation are available in the text and appendix. There are 306 HRR in the US.

tection Agency Air Quality Index (2012-2016), Census of Governments (2012), National Center of Education Statistics, American Community Survey (2012-2016). Refer to section 6.1 for a more complete discussion on the amenity index estimation.

4.2 New Dataset: Physician Work History Panel

To understand when the physician chooses their first job, I need to know when they finish their training. To do so, I use two more datasets. First, I get all the data from Medicare Physician Compare. Since it is a file created by Medicare, I can match physicians through the NPI. This directory allows me to see their graduation year as well as their medical school for about half of the physicians. Unfortunately, this dataset seems to not report the medical school name for many physicians, by absorbing all of them under “Other.” To resolve this issue and be able to account for their residency training as well, I scrape data from the internet directories, in particular Doximity.com, where physicians publish a wealth of information regarding their training and affiliations. I am able to then have a sample of physicians scraped off the internet matched not only to their specialty and office address, but also to their graduation year, medical school, residency, internship and fellowship participation, among other characteristics. From their graduation year, I add the number of years of residency depending on their specialty, as well as an extra year each if they do a year of internship or fellowship. Having done this, I know the year they pick their first job, and I then select out those that pick their first job between 2012 and 2016, to be able to match the variables in the Medicare data. I call the full panel I created the “Physician Work History Panel.”

A selected sample of variables from the Physician Work History Panel is available in Table 1 to provide the Reader with a flavor of the collected data.

To have a proxy of the level of quality/skill that physicians could be characterized by, I use the US News rankings for medical schools, by specialty. Moreover, I use the Doximity Residency Navigator ranking to proxy for the quality of the residency program, by specialty. Recall that each physician declares their specialty and I am able to control for multiple ones if they identify with more than one. This allows me to identify a normalized ranking of their medical school and residency, as well as to know whether they have attended a top-10, top-30, top-50 school and residency for their specific specialty, as well as whether or not they attended a school or residency in the bottom quarter percentile of the rankings. Moreover, even though I do not observe physicians' immigrant status, I can get insight from the medical school they attended. This is interesting because foreigners can obtain a visa waiver if they practice in health professional shortage areas or medically underserved areas. I match the medical schools I observe to their addresses and mark those outside the United States, assuming that a person that studied abroad for medical school is, in fact, foreign. The foreign parameter would be, if anything, understated, since I cannot observe immigrants that migrated prior to medical school, but it is a good-enough proxy for the analysis at hand. By including this in the analysis, I am able to see whether there is, in fact, a higher choice of rural placements by foreigners to be able to take advantage of visa benefits.

The other crucial variable for the analysis is the salary received by physicians. For this, I utilize the Occupation Employment Statistics (OES) at the Metropolitan Statistical Area (MSA)-level for the main occupational divide (primary care physicians vs. surgeons) and subspecialties defined by BLS. The fixed salary data gives an idea of the incentives that physicians face from a pure salary perspective, but they do not include the procedure revenue dependent on fee-for-service income. As shown in Falcettoni (2018), this procedure revenue highly depends on the procedures carried out, but the mix of procedures carried out on average depends not only on the specialty of the physician, but on the location where the physician is located as well. For this reason, I believe that it is important to consider both income sources in the analysis. More details on the income composition can be found in Section 3.

The choice of a location that displays a shortage of physicians is often associated with high benefits in terms of loan forgiveness. Since my entire dataset is post-Obamacare (when these benefits became higher and tax-deductible), the control for those areas designated as health professional shortage areas captures the loan forgiveness effect. To capture the benefit given by this particular element, I match the medical schools to their tuitions and calculate the amount of student debt that the resident would be facing. The medical debt then disappears as an expense if they were to choose areas offering loan forgiveness.

Finally, I am of course aware that not every physician would necessarily be online and captured in my dataset. Naturally, the fact that I focus on the first job out of residency starting in 2012 is very useful, as these physicians are younger and more likely to take advantage of internet directories. Regardless, I use more-aggregated public data sources to see how representative the dataset I created is. The sample collected can explain almost the entirety of the variance in the data, with an R^2 of 93.74 percent for primary care physicians and 96.47 percent for specialists if I regress the county-by-county data of the sample I collected against the public county-by-county data released by the Area Health Resource File. Notice that the sample should not display a perfect fit. As a matter of fact, since the data is based on the Area Resource File data, residents and non-degree primary care physicians are included in the estimates, while my sample excludes them. A similarly-high representativeness is true also for the stock of physicians on a county-by-county basis. I also perform some validation exercise on the whole population of primary care physicians and specialists by year, by looking at the Area Health Resource File by specialty and year. Table 2 reports a sample of the results.[3]Overall, the sample considered matches at least 85 percent of the total physician population

[3] Further statistics on the fit such as the county-by-county fit discussed above are available upon request.

Table 2: Physicians Work History Panel vs. Leading Data Source

	Physician Work History Panel	Leading Data Source	percent
Total Year 3 Residents	31,759	37,617 (Board)	84 percent
2012 Primary Care Population	198,310	237,346 (AHRF)	84 percent
2013 Primary Care Population	201,527	242,955 (AHRF)	83 percent
2014 Primary Care Population	203,445	244,638 (AHRF)	83 percent
2015 Primary Care Population	205,064	245,983 (AHRF)	83 percent
2016 Primary Care Population	186,312	247,069 (AHRF)	75 percent

Notes: The Physician Work History Panel is a dataset that I created and that provides physician-level data on their training and on the work they currently carry out. The Panel is created through two main data sources: first, I use the Medicare Part B Utilization and Payment data; second, I scrape physician directories (mainly Doximity.com) to be able to determine their medical training (medical school, , residency). I then use the medical school information to infer the level of student debt they would be facing and I collect Bureau of Labor Statistics data on wages by occupation title to collect information on salaries. While this panel covers almost the universe of physicians, I focus on residents that finish residency and enter the medical job market in this paper. This table provides one of many validation exercises to compare the Physician Work History Panel to commonly used data sources for similar statistics. Since I do not have access to any dataset that measures the number of residents on the medical job market by year, I approximate it by the number of Year 3 Residents. Notice that some Year 3 Residents continue their training. Therefore, the leading data source should provide a slightly higher value than the number of residents in my panel, as it is indeed true. Next, I also show in this table how the Physician Work History Panel compares in terms of primary care physician population. To do so, I utilize the Area Health Resource Files by county and I aggregate my data up to the county level to enable a comparison. The Physician Work History Panel is able to reproduce the levels and changes in the physician population closely. Similar statistics are true for specialists as well. More details on the data collection and sample validation are available in the text and appendix.

at a national level. These facts are particularly key because the representativeness of the sample at the geographical level is crucial for the analysis to be qualitatively correct and to be able to evaluate current and alternative policies aimed at bringing physicians to more-rural areas.

Finally, I also check the retention rate in the public data vs. in the dataset I collected. The data I collected produces a within-state retention of 51.4 percent overall, while public data reports a 54 percent state retention nationally, further confirming the sample fit.

5 Model

From now on, for ease of exposition, I will refer to location choice, city choice, and hospital referral region (HRR) choice interchangeably. The goal of HRRs is to define areas that are self-contained markets for physicians, so that the majority of patients living in that area would be able to remain within their HRR for any visit they need. There are 306 HRRs in the United States. The model shown here is written as the model of a physician picking their location choice. The model is generalized in the obvious way in the estimation when I focus on first-job choices only, recalling that the change in the total population of physicians of either type is given by the population of new physicians of either type.

5.1 Physician Supply (i.e. the Demand for Locations)

I set up physicians' choice of a location as a structural static discrete choice location. Physicians pick a hospital referral region (HRR) to live in. The goal of HRRs is to define areas that are self-contained markets for primary care, so that the majority of patients living in that area go to primary care physicians within that area. Since physicians are picking the location of their job, HRRs provide me with the area that will constitute their market for medical procedures they will carry out. The outside choice is set to be equal to the HRR named "La Crosse, WI"[4] every year.

Physicians are grouped into two main specialty groups: primary care and specialty care. Individuals i are heterogeneous along their specialty k and two demographic characteristics ℓ: the quality of the residency they completed as a proxy of their skill, q_i, and whether they are foreign, f_i. The quality of the residency attended is used as a proxy of the skill level of the physicians.

To approximate for the fact that higher-skilled residents have more options available to them than their lower-skilled counterparts, I rank both jobs and individuals by their ranking. I then approximate the quality of a location by the average quality of the jobs within that same location. The top 1 percent of physicians each year is the only group who has access to the locations containing the top-1 percent jobs. Of course, they are also able to pick any location with jobs that are lower-ranked than they are. The top 10 percent of physicians can pick any location that is equal or lower than their ranking. From the top 50 percent onward, physicians can pick locations that are immediately above their ranking or below theirs. For example, a top 40 percent individual has all locations with an average quality of jobs that are below the 30 percent threshold in her choice set. Importantly, this quality ladder really only matters at the bottom, when jobs are exhausted and the better locations, which often might have both higher amenities and higher pay, are already filled.

Cities do not only differ by the wages and the physician-type mix. They differ by the level of amenities. I collect amenities on a variety of characteristics, grouped into seven main categories: cost of housing, entertainment, safety, transportation, education, crime, and environment. Amenities x_{jt} are treated as exogenous in this setting (physicians are one occupation only that will not influence the amenities in that location).

Finally, I allow for physicians' preference for locations that are close to where they completed their residency. As discussed, data evidence shows this preference to be strong and an important factor in physicians' choices. Therefore, preferences of workers with the same demographic characteristics ℓ for a HRR j can differ due to preferences to remain within the residency's HRR and state. Their preferences can also differ through the income because because of the individual-specific student loans.

As commonly done in the literature, I express physicians' preferences as the indirect utility function physicians receive when picking HRR j in year t. I suppressed the time index t for ease of read. New physicians pick a location within the whole nation, but they compete with the graduates that are also picking a location in the same year. Recall that residents are differentiated along their specialty group, their quality, and their foreign status.

I utilize the micro-data that I collected to let physicians differ in how they value the net income offered in different locations. The endogeneity issue that is commonly present within the mean utility parameters will then disappear, as compensation will not be contained in the mean utility anymore and there are no restrictions imposed between the unobserved amenities and the error term.

I let physicians differ in their preferences not only due to the location of their residency and the idiosyncratic shock, but also due to the net income they receive. Since I know the medical school physicians attended and I use this information to calculate the average payment of student loans, physicians actually differ in the net

[4]This HRR contains Minneapolis, MN, so it was picked as tribute to the University of Minnesota. Go Gophers!

income they would receive in the same location. Therefore, the specification I run is the following:

$$\max_j u_{ij} = \overbrace{\beta^{k,\ell} x_j + \xi_j^{k,\ell}}^{\delta_j^{k,\ell}} + \overbrace{\alpha^{k,\ell} y_{ij} + \beta_j^{k,\ell} x_{ij}}^{\mu_{ij}} + \epsilon_{ij} \tag{2}$$

where

$$u_{ij} = \begin{cases} \delta_j^{k,\ell} & \text{mean utility} \\ +\mu_{ij} & \text{stochastic coefficients} \\ +\epsilon_{ij} & \text{iid T1EV error term} \end{cases} \tag{3}$$

where k represents the specialty group of the physician (primary care vs. specialty care) and ℓ represents the demographics of the physicians, namely their quality ranking q_i and their foreign status f_i. x_j are the location characteristics (the observed amenities), $\xi_j^{k,l}$ are location-year unobservables. y_{ij} is net income, which includes, as mentioned before: the expected wage income, the expected procedure revenue, the expected expenses (housing costs, malpractice insurance, student loans), and the expected subsidies (loan forgiveness affects the amount of student loans, while salary incentives increase the expected wage income). The ϵ_{ij} are drawn from Type 1 extreme value distribution and are independent and identically distributed across physicians, locations, and years.

Then, the probability that a physician i of type ℓ in specialty k picks location j in a given year is:

$$\hat{s}_{ijt} = \frac{\exp\left\{\delta_{jt}^{k,\ell} + \mu_{ijt}\right\}}{\sum_{m=1}^{M} \exp\left\{\delta_{mt}^{k,\ell} + \mu_{imt}\right\}} \tag{4}$$

where M is equal to the number of HRRs.

The overall portion of physicians in specialty k picking location j in a given year across all demographic characteristics can be found by summing the individual market shares across the individuals within each type and across types. The number of primary care and specialty care physicians in j at time t are respectively:

$$N_{jt}^{PC} = \sum_{\ell \in q,f} \sum_{i=1}^{N_{\ell t}^{PC}} \frac{\exp\left\{\delta_{jt}^{PC,\ell} + \mu_{ijt}\right\}}{\sum_{m=1}^{M} \exp\left\{\delta_{mt}^{PC,\ell} + \mu_{imt}\right\}} N_{\ell t}^{PC} \tag{5}$$

$$N_{jt}^{SP} = \sum_{\ell \in q,f} \sum_{i=1}^{N_{\ell t}^{SP}} \frac{\exp\left\{\delta_{jt}^{SP,\ell} + \mu_{ijt}\right\}}{\sum_{m=1}^{M} \exp\left\{\delta_{mt}^{SP,\ell} + \mu_{imt}\right\}} N_{\ell t}^{SP} \tag{6}$$

5.1.1 Net income in the mean utility $\delta_j^{k,\ell}$

The discrete-choice literature usually includes the monetary value in the mean utility. This is usually due the fact that the monetary value does not differ by the individual i, as it is the case, for example, with cereal, where the price does not differ by the purchaser. This same assumption has usually been kept in analyses where the monetary value indeed varies for each individual, as it is the case here. To compare what the results would look like under an average-income scenario, I use the micro data to calculate an average net income that would be available in a given location j for physicians of type ℓ in specialty k and estimate the model again, using the instrumental variables that will be discussed below to jointly estimate the supply and demand coefficient of compensation in this scenario. Note that this average net income still includes all the elements present in the micro data (salary, procedure revenue, housing cost, malpractice insurance, student loans), but it simply represents what the average net income is for a physician of type ℓ in specialty k in location j. Section A.2 in the Appendix walks through the setup of the model, the estimation, and the results from the paper in this alternative specification. I show in the results that the coefficients found do not differ qualitatively from those found letting net income vary for each individual. The correct measure of net income and the correct specification of the choice sets are the key to the results found in this paper. However, quantitatively, the extent to which physicians respond to incentives of course depends on the incentive that they would individually receive, so that the parameters vary between the two specifications because of this.

5.1.2 The Independence of Irrelevant Alternatives (IIA)

Notice that the IIA property does not hold here. First, the individual preferences μ_{ij} allow for correlation in preferences within areas and within states due to physicians' preferences to have a preference to remain close to their residency. Within each area, preferences vary by specialty and by ℓ, therefore breaking the IIA property within each region.

5.2 Physician Demand

I model health firms in the city to use capital as well as a composition of primary care and specialty care workers, so that the composition of the two types in a city matters for production. The reasoning for this is to imagine physicians as employed by clinics/hospitals/physician offices, which use machinery as capital and a given mix of primary care and specialty care physicians to produce a given health good, which I assume, for now, to be one and homogenous in production across the two types.

Following the results discussed in Falcettoni (2018), which show that the mix of procedures carried out depends on the types of physicians in the same location, I allow for the productivity of primary care and specialty care physicians to depend on the mix between the two types. Physicians then receive a wage as well as procedure revenue, the sum of which determines the total compensation. Procedure revenue, and therefore compensation, depends on the specific procedures carried out. The wage endogenously responds to changes in physician employment, while procedure revenue endogenously responds to the mix of types of physicians in the city. The wage is also impacted by the changes in the physician's productivity caused in equilibrium by the ratio between primary and specialty care physicians. Therefore, total compensation clearly responds to changes in the employment of either type of physicians. Physicians could also be self-employed. This set-up can easily include self-employment by imagining the physician to be simply employed in her own firm. Since I observe whether or not physicians are in an office or in a facility setting, I can allocate income accordingly.

Each HRR j produces a medical good through a high number of homogenous firms (index suppressed) that employ primary care physicians (N_{jt}^{PC}) and specialists (N_{jt}^{SP}) and use machinery (capital K_{jt}). Production follows a Cobb-Douglas function where the total number of physicians employed (D_{jt}) follows a CES function:

$$\begin{aligned} M_{jt} &= D_{jt}^{\alpha} K_{jt}^{1-\alpha} && (7) \\ D_{jt} &= \left(\theta_{jt}^{PC} \left(N_{jt}^{PC} \right)^{\rho} + \theta_{jt}^{SP} \left(N_{jt}^{SP} \right)^{\rho} \right)^{\frac{1}{\rho}} && (8) \\ \theta_{jt}^{PC} &= f_{PC} \left(N_{jt}^{PC}, N_{jt}^{SP} \right) \exp \left(\epsilon_{jt}^{PC} \right) && (9) \\ \theta_{jt}^{SP} &= f_{SP} \left(N_{jt}^{PC}, N_{jt}^{SP} \right) \exp \left(\epsilon_{jt}^{SP} \right) && (10) \end{aligned}$$

which leads to an elasticity of substitution between primary care physicians and specialists equal to $\frac{1}{1-\rho}$. This setup for physician demand is borrowed from the extensive literature on the wage differential between high- and low-skilled workers, such as Katz and Murphy (1992) and, more recently, Diamond (2015).

$\theta_{jt}^{k}\ \forall k = PC, SP$ allow for primary care physicians and specialists to have two different productivities. The two error terms are the exogenous factors that affect the productivity levels, while the function that takes into consideration the employment of either type of physicians endogenously affects the productivity of either type. This setup works well in my model since primary care physicians behave differently and can carry out different procedures when specialists are not around. Therefore, their productivity highly depends on the number of physicians of either type within a location. The two functions are not specified by a parametric form to not make strong assumptions on the way that the number of physicians directly impacts productivity.

Since the firms are perfectly competitive, they hire until total compensation is equal to the marginal product of labor. The capital market is assumed to be frictionless and national, so that firms can get capital at price p_t, equal across all locations. Finally, since firms are homogenous and the production function is Cobb-Douglas, the firm problem is representative to each location's physician demand. Solving the above problem and log-linearizing:

$$\begin{aligned} \text{comp}_{jt}^{PC} &= a_t + (1-\rho)\, d_{jt} + (\rho - 1)\, n_{jt}^{PC} + \log \left(f_{PC} \left(N_{jt}^{PC}, N_{jt}^{SP} \right) \right) + \epsilon_{jt}^{PC} && (11) \\ \text{comp}_{jt}^{SP} &= a_t + (1-\rho)\, d_{jt} + (\rho - 1)\, n_{jt}^{SP} + \log \left(f_{SP} \left(N_{jt}^{PC}, N_{jt}^{SP} \right) \right) + \epsilon_{jt}^{SP} && (12) \end{aligned}$$

$$D_{jt} = \left(f_{PC} \left(N_{jt}^{PC}, N_{jt}^{SP} \right) \exp \left(\epsilon_{jt}^{PC} \right) \left(N_{jt}^{PC} \right)^{\rho} + f_{SP} \left(N_{jt}^{PC}, N_{jt}^{SP} \right) \exp \left(\epsilon_{jt}^{SP} \right) \left(N_{jt}^{SP} \right)^{\rho} \right)^{\frac{1}{\rho}} \quad (13)$$

where lowercase letters stand for log-variables, comp is the logarithm of the total compensation physicians receive (salary and procedure revenue) and a_t is a constant, given by: $a_t = \log \left(\alpha \left(\frac{(1-\alpha)}{p_t} \right)^{\frac{1-\alpha}{\alpha}} \right)$.

The physician demand equations can be approximated with log-linear aggregate physician demand, as commonly done in the labor literature:

$$\begin{aligned} \text{comp}_{jt}^{PC} &= \beta_{0,pc} + \gamma_{pc}^{pc} n_{jt}^{PC} + \gamma_{sp}^{pc} n_{jt}^{SP} + \epsilon_{jt}^{PC} && (14) \\ \text{comp}_{jt}^{SP} &= \beta_{0,sp} + \gamma_{pc}^{sp} n_{jt}^{PC} + \gamma_{sp}^{sp} n_{jt}^{SP} + \epsilon_{jt}^{SP} && (15) \end{aligned}$$

where the total compensation and employment in each j are data, while the errors are unobserved, and γ_{pc}^{pc}, γ_{sp}^{pc}, γ_{pc}^{sp}, γ_{sp}^{sp} are the parameters to be estimated.

5.3 Equilibrium

The equilibrium is given by a set of total compensation (the sum of average salary and procedure revenue by specialty and location) and quantity of physicians for every location j in every year t, $\left(\text{comp}_{jt}^{PC*},\ \text{comp}_{jt}^{SP*},\ N_{jt}^{PC*},\ N_{jt}^{SP*}\right)_{\forall j,t}$, such that:

1. Demand for primary care physicians equals supply of primary care physicians in each city:

$$\begin{cases} N_{jt}^{PC^*} & = \sum_{\ell \in q,f} \sum_{i=1}^{N_{\ell t}^{PC}} \frac{\exp\left\{\delta_{jt}^{PC,\,\ell} + \mu_{ijt}\right\}}{\sum_{m=1}^{M} \exp\left\{\delta_{mt}^{PC,\,\ell} + \mu_{imt}\right\}} N_{\ell t}^{PC} \\ \text{comp}_{jt}^{PC} & = \beta_{0,pc} + \gamma_{pc}^{pc} n_{jt}^{PC^*} + \gamma_{sp}^{pc} n_{jt}^{SP} + \epsilon_{jt}^{PC} \\ n_{jt}^{PC^*} & = \log N_{jt}^{PC^*} \end{cases} \tag{16}$$

2. Demand for specialists equals supply of specialists in each city:

$$\begin{cases} N_{jt}^{SP^*} & = \sum_{\ell \in q,f} \sum_{i=1}^{N_{\ell t}^{SP}} \frac{\exp\left\{\delta_{jt}^{SP,\,\ell} + \mu_{ijt}\right\}}{\sum_{m=1}^{M} \exp\left\{\delta_{mt}^{SP,\,\ell} + \mu_{imt}\right\}} N_{\ell t}^{SP} \\ \text{comp}_{jt}^{SP} & = \beta_{0,sp} + \gamma_{pc}^{sp} n_{jt}^{PC} + \gamma_{sp}^{sp} n_{jt}^{SP^*} + \epsilon_{jt}^{SP} \\ n_{jt}^{SP^*} & = \log N_{jt}^{SP^*} \end{cases} \tag{17}$$

3. The compensation clears the market:

$$\begin{cases} \text{comp}_{jt}^{PC*} & = \beta_{0,pc} + \gamma_{pc}^{pc} n_{jt}^{PC^*} + \gamma_{sp}^{pc} n_{jt}^{SP*} + \epsilon_{jt}^{PC} \\ \text{comp}_{jt}^{SP*} & = \beta_{0,sp} + \gamma_{pc}^{sp} n_{jt}^{PC*} + \gamma_{sp}^{sp} n_{jt}^{SP^*} + \epsilon_{jt}^{SP} \end{cases} \tag{18}$$

Recall that total compensation is contained in the physician supply equation as part of the net income received by physicians.

6 Estimation

6.1 Amenity Index

The computation of the amenity index follows Diamond (2015). The amenity index is meant to capture the heterogeneity in the amenity bundle available in every location. To approximate for such bundle as closely as possible, I include eight different categories: retail, education, environment, health, crime, transportation, long commute, and traffic. Recall that job availability and the characteristics of the patients available are analyzed separately from this index.

Retail includes the number of clothing stores per capita calculated using the U.S. Census data on apparel stores; education includes the county spending per pupil up to secondary school as well as state spending per capita on libraries, primary, and secondary schools; environment includes the investment in parks and green spaces, the number of parks, the number of days marked with pollution, the median level of pollution,

Table 3: Amenity Index

Variables	First Components, Primary Care	First Components, Specialists
Clothing Stores	0.0525	0.0578
Education	0.4828	0.4764
Environment	0.3529	0.3800
Health	0.4595	0.4716
Crime	-0.0201	-0.0273
Transportation	0.4550	0.4620
Long Commute to Work	-0.3257	-0.3038
Traffic	-0.3392	-0.3207

Notes: These results come from the estimation of the amenity index discussed in the paper. Each value represents the first components obtained through principal component analysis (using correlation). First, I run principal component analysis on the single categories. Then I run principal component analysis again on the single categories to obtain the first components shown in this table, representing the weight of each category on the final index. More details are available in the text. The first components of each category are available in the online appendix.

the number of good, moderate, and unhealthy days as measured by the Environmental Protection Agency (using the air quality index); the health facility index includes investment in hospitals and health facilities; crime includes the number of correction facilities, violent crimes, murders, rapes, robberies, aggravated assaults, property crimes, burglaries, thefts, and motor thefts per capita; transportation facilities includes highways, airports, parkings, and harbors per capita; commute includes the percentage of people within the area commuting, by length of commute, from up to 14 minutes to over an hour; finally, traffic includes the percentage of people that commute by car, by length of commute, to proxy for the number of cars that would be in an area. I use principal component analysis (PCA) based on correlation (because of the different scales of the variables) to build the index. I extract the first component for each of the eight categories first, and I then run PCA again on the individual categories mentioned above to generate the full amenity index.

Within each individual index, the sign of the weight assigned to each categories correctly puts the right weight on the single components. Environment puts negative weight on the days with pollution, the median level of pollution, moderate, and unhealthy days, correctly picking up that those factors decrease the quality of the air, while all other factors improve it. Long commute puts a negative weight on short commutes up to 24 minutes, with factor loadings increasing in the time of commute, correctly picking up that short commutes decrease the commute length, and more so, the shorter the commute. Traffic puts a negative weight on short commutes by car below 20 minutes, correctly picking up that areas in which people commute little by car to get to work are characterized by less traffic. All the other indices put positive weights on the different factors, correctly picking up that each of them positively contributes to the individual index itself.

I finally run PCA again on the single indices to create the amenity index used in the paper. The index for both specialty groups is presented in Table 5. The index is able to capture that crime, a long commute, and traffic are negative attributes of an area. It is also able to capture that, instead, a high number of stores, a high quality of the environment, and a high investment in education, health facilities, and transportation are all positive attributes of an area. To check that this is in fact correct, I rank the hospital referral regions according to the index level. The amenity index ranked HRRs inside and around New York, Chicago, DC, San Francisco, and Seattle all at the top of the list. Those cities are all places generally considered to have a high level of amenity, reinforcing the validity of the index.

6.2 Instrumental Variables

I use a variation of the Bartik (1991) instruments to be able to identify my parameters of interest. Bartik shocks are generally defined as local labor demand shocks driven by the share of the city's employment in that industry with respect to the importance of that same industry nationally. Bartik instruments are therefore able to measure the change in a region's labor demand that is induced by changes in the national demand for different industries' products.

Following this logic, I utilize my micro data on physicians' transactions at the procedure level for Medicare to be able to identify labor demand shocks that are uncorrelated with labor supply shocks. The identifying assumption in this setting is that labor demand is procedure-specific, while labor supply is not. In other words, for example, patients care about having someone carrying out an EKG, if they do not have somebody providing it already, more than X-Rays, if someone already carries those out, but the physician herself, already assigned to a specialty, does not have a procedure-specific taste.

In order for this to hold, the rates I consider have to be exogenous, otherwise other issues of endogeneity could be present. The reason I use Medicare fee-for-service reimbursement rates is that these procedure rates are set by policy. Reimbursement rates change for the whole nation according to the decisions made by the reimbursement committee. This allows me to evaluate the monetary impact of the productivity shocks on procedures through the shock on the procedure revenue:

$$B_{jt}^{SP} = \sum_{q \in treatments} \frac{N_{j,t-k}^{SP,\,q}}{\sum_{q'} N_{j,t-k}^{SP,\,q'}} \log\left(\frac{\sum_{m\neq j} ProcRev_{m,t}^{SP,q}}{\sum_{m\neq j} ProcRev_{m,t-k}^{SP,q}}\right) \tag{19}$$

$$B_{jt}^{PC} = \sum_{q \in treatments} \frac{N_{j,t-k}^{PC,\,q}}{\sum_{q'} N_{j,t-k}^{PC,\,q'}} \log\left(\frac{\sum_{m\neq j} ProcRev_{m,t}^{PC,q}}{\sum_{m\neq j} ProcRev_{m,t-k}^{PC,q}}\right) \tag{20}$$

which can be written in an equivalent way as:

$$B_{jt}^{SP} = \sum_{q \in treatments} \frac{N_{j,t-k}^{SP,\,q}}{N_{j,t-k}^{SP}} \left(\overline{ProcRev}_{m\neq j,t}^{SP,q} - \overline{ProcRev}_{m\neq j,t-k}^{SP,q}\right) \tag{21}$$

$$B_{jt}^{PC} = \sum_{q \in treatments} \frac{N_{j,t-k}^{PC,\,q}}{N_{j,t-k}^{PC}} \left(\overline{ProcRev}_{m\neq j,t}^{PC,q} - \overline{ProcRev}_{m\neq j,t-k}^{PC,q}\right) \tag{22}$$

where $\frac{N_{j,t-k}^{SP,\,q}}{N_{j,t-k}^{SP}}$ is the share of specialists carrying out procedure q at the beginning of the period in HRR i and $\frac{N_{j,t-k}^{PC,\,q}}{N_{j,t-k}^{PC}}$ is, equivalently, the share of primary care physicians carrying out procedure q in 2012 in HRR i, while $\left(\overline{ProcRev}_{m\neq j,t}^{SP,q} - \overline{ProcRev}_{m\neq j,t-k}^{SP,q}\right)$ is the growth in average reimbursement to specialists for procedure q between 2012 and 2016 in all other HRRs but HRR j, and $\left(\overline{ProcRev}_{m\neq j,t}^{PC,q} - \overline{ProcRev}_{m\neq j,t-k}^{PC,q}\right)$ is the equivalent measure for primary care physicians.

6.2.1 Instrument validity

This instrument will be valid, relevant, and exogenous under the following restrictions:

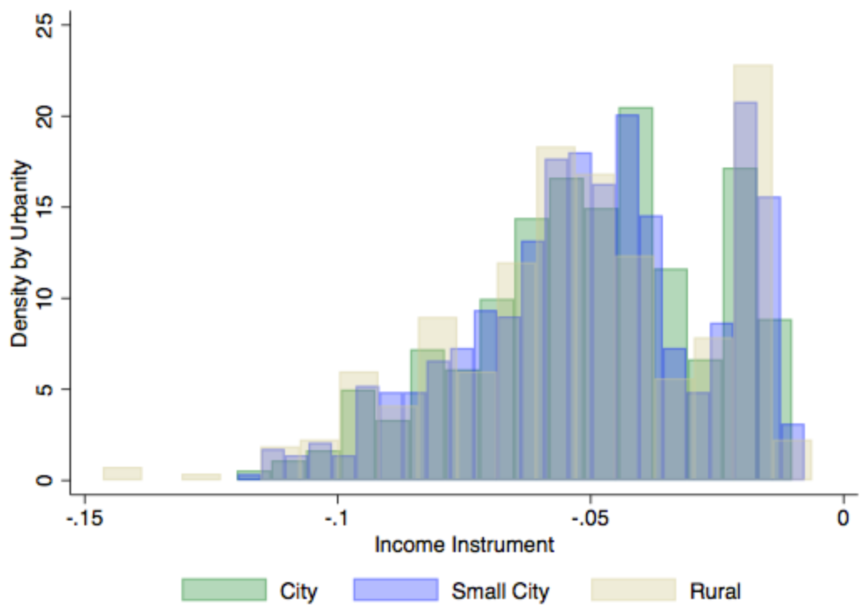

Figure 2: Income Instrument Density by Urbanity Index

Notes: This figure shows density distribution of the income instrument created, across and within urbanity types. Source: CMS.

1. There needs to be enough variation in the national growth rates across different procedures (i.e. $\left(\overline{ProcRev}^{SP,q}_{m\neq j,t} - \overline{ProcRev}^{SP,q}_{m\neq j,t-k}\right)$ and $\left(\overline{ProcRev}^{PC,q}_{m\neq j,t} - \overline{ProcRev}^{PC,q}_{m\neq j,t-k}\right)$ are different for different q).

   This is easily respected because rate changes are set by policy and the reimbursement rates are different for different procedures, so there are enough changes across the five years considered for this to hold.

2. There needs to be enough variation in the share of physicians carrying out procedure q across areas j ($\frac{N^{SP,\,q}_{j,t-k}}{N^{SP}_{j,t-k}}$, $\frac{N^{PC,\,q}_{j,t-k}}{N^{PC}_{j,t-k}}$ are different for different j), or in other words, the mix of procedures carried out has to vary by location. This fact is evident in the data as it is briefly discussed in Section A.1 in the Appendix.

3. The exclusion restrictions have to be respected:

   (a) No procedure q is concentrated in one area only.

   (b) There is no supply effect or shock that drives the Bartik instrument.

To test the third point, I need to test that there is enough variation in the procedure mix carried out in different places and that these differences in the procedure mix are uncorrelated with location unobservables. First, as shown in Figure 2, the instrument created exhibits high variation across and within the urbanity level.

Second, the exogeneity of the instrument could be challenged thinking that differences in the procedure mix carried out by physicians were correlated with the same characteristics that make people decide to live in different locations. This is exactly why I exploit the changes in the policy reimbursement rates and the procedure count observed across the years in my data. Since the instrument is calculated using the changes in procedure revenue observed between 2012 and 2016, it is unlikely that unobserved location amenities are systematically correlated to these changes across time. Moreover, notice that the shock on procedure revenue over time always excludes the area for which the instrument is built to eliminate local effects. Even though I run a two-step simultaneous GMM, I run a 2SLS estimation to get a feeling of the first-stage regression results and determine the instrument strength. The F-Stats are high (~68), suggesting that the instrument performs well in income prediction. The full estimation passes the over-identification Hansen's J-test as well, with a p-value equal to 0.4182 in the first case and 0.4049 in the other.[5]

As previously mentioned, Bartik demand shocks are uncorrelated to local city characteristics since they are built based on national shocks to the physicians' procedure revenue level. Finally, I interact the Bartik instruments with malpractice insurance premia. Notice that changes in the malpractice insurance premia do not depend on the attractiveness of the city, but vary substantially within and across locations. The interaction with malpractice insurance premia acts as a tiebreaker: if two places experience the same shock and are equivalent along every other dimension, the place that experiences the higher increase in malpractice insurance becomes less attractive. Simultaneous estimation guarantees identification thanks to the instruments' ability to affect all endogenous variables at the same time and simultaneously identify the parameters of interest.

6.3 Physician Supply

The estimation of physicians' preferences is a two-step estimation procedure, following BLP (1995). First, I use maximum likelihood to identify how much physicians of either type in each year want to live in each location, obtaining the mean utility level for each location $\left(\delta_j^{k,\ell}\right)$ and the individual coefficients (contained in μ_{ij}). Recall that in the estimation, $\beta_j^{k,\ell}$ is a vector of two parameters, representing the semi-elasticity of demand with respect to whether the location chosen is within the same state $\left(\beta_{state}^{k,\ell}\right)$ and area $\left(\beta_{HRR}^{k,\ell}\right)$ of residency. I then use these and $\alpha^{k,\ell}$ in the second step to estimate the trade-off physicians face between wages and other characteristics in their location choice through simultaneous equation non-linear GMM, using moments on physicians' preferences and physician's demand. Standard errors are clustered by hospital referral region.

Let $\mathbb{I}_{ij}$ be a dummy variable that takes value one if physician i chooses location j. Recall that the probability that a physician i of type ℓ in specialty k picks location j in a given year is:

$$\hat{s}_{ijt}^{k,\ell} = \frac{\exp\left\{\delta_{jt}^{k,\ell} + \alpha^{k,\ell} y_{ijt} + \beta_j^{k,\ell} x_{ijt} + \epsilon_{ijt}\right\}}{\sum_{m=1}^{M} \exp\left\{\delta_{mt}^{k,\ell} + \alpha^{k,\ell} y_{imt} + \beta_m^{k,\ell} x_{imt} + \epsilon_{imt}\right\}} \tag{23}$$

[5] The rule of thumb is that a high p-value indicates a good model fit, and a p-value that is not too high suggests that there are no other issues related to the model.

Recall that the same apparent "compensation offer", represented by y_{ijt} in a location is perceived differently by different individuals, as they each carry their level of student loans and expenses. By utilizing the micro-data, the net income variable is calculated for each individual. Through this methodology, individuals of the same type ℓ are still allowed to differ in their preferences for locations according to the total net income that they individually would face in each location. The log-likelihood function is then given by:

$$\mathcal{L}^{k,\ell}\left(\boldsymbol{\delta_{jt}^{k,\ell}},\, \alpha^{k,\ell},\, \boldsymbol{\beta_j^{k,\ell}}\right) = \frac{1}{N_\ell^k}\sum_{i=1}^{N_\ell^k}\sum_{j=1}^{J}\mathbb{I}_{ij}\log\left(\hat{s}_{ijt}^{k,\ell}\left(\boldsymbol{\delta_{jt}^{k,\ell}},\, \alpha^{\mathbf{k},\ell},\, \boldsymbol{\beta_j^{k,\ell}};y_{ij},\, x_{ij}\right)\right) \tag{24}$$

This log-likelihood cannot be estimated directly because it would require a search over $J+1$ parameters (which is problematic, since J>300). Therefore, I instead use the BLP (1995) contraction that finds the parameters of interest by matching the observed shares in the data $s_{jt}^{k,\ell}$ to the estimated shares $\hat{s}_{ijt}^{k,\ell}\left(\boldsymbol{\delta_t^{k,\ell}},\, \alpha^{\mathbf{k},\ell},\, \boldsymbol{\beta_j^{k,\ell}}\right)$:

$$T\left(\delta_{\mathbf{j}}^{\mathbf{k},\ell}\right) = \delta_j^{k,\ell} + [\log\left(s_j^{k,\ell}\right) - \log\left(\hat{s}_j^{k,\ell}\left(\boldsymbol{\delta_j^{k,\ell}},\, \alpha^{k,\ell}\,\boldsymbol{\beta_j^{k,\ell}}\right)\right) \tag{25}$$

which reduces the estimation considerably, since for every $\boldsymbol{\beta_j^{k,\ell}}$, there exists a unique $\delta_j^{k,\ell} = \left(\delta_1^{k,\ell}, ..., \delta_J^{k,\ell}\right)$ that matches the observed and estimated shares. The log-likelihood is then given by:

$$\mathcal{L}^{k,\ell}\left(\boldsymbol{\delta^{k,\ell}}\left(\boldsymbol{\beta_j^{k,\ell}}\right),\, \alpha^{k,\ell},\, \boldsymbol{\beta_j^{k,\ell}}\right) = \frac{1}{N_\ell^k}\sum_{i=1}^{N_\ell^k}\sum_{j=1}^{J}\mathbb{I}_{ij}\log\left(\hat{s}_{ijt}^{k,\ell}\left(\boldsymbol{\delta^{k,\ell}}\left(\boldsymbol{\beta_j^{k,\ell}}\right),\, \alpha^{k,\ell},\, \boldsymbol{\beta_j^{k,\ell}};y_{ij},\, x_{ij}\right)\right) \tag{26}$$

The recovered $\boldsymbol{\delta_j^{k,\ell}}$ are now simply given by the amenities of a location, i.e. the common component of the utility that individuals within each type agree upon. The second stage finally allows me to recover $\xi_{jt}^{k,\ell}$.

6.4 Physician Demand

Recall the labor demand equations:

$$\text{comp}_{jt}^{PC} = \beta_{0,pc} + \gamma_{pc}^{pc} n_{jt}^{PC} + \gamma_{sp}^{pc} n_{jt}^{SP} + \epsilon_{jt}^{PC} \tag{27}$$

$$\text{comp}_{jt}^{SP} = \beta_{0,sp} + \gamma_{pc}^{sp} n_{jt}^{PC} + \gamma_{sp}^{sp} n_{jt}^{SP} + \epsilon_{jt}^{SP} \tag{28}$$

and take the difference of the variables from the base values in 2012, utilizing the panel dimension of the data:

$$\Delta\text{comp}_{jt}^{PC} = \gamma_{pc}^{pc}\Delta n_{jt}^{PC} + \gamma_{sp}^{pc}\Delta n_{jt}^{SP} + \Delta\epsilon_{jt}^{PC} \tag{29}$$

$$\Delta\text{comp}_{jt}^{SP} = \gamma_{pc}^{sp}\Delta n_{jt}^{PC} + \gamma_{sp}^{sp}\Delta n_{jt}^{SP} + \Delta\epsilon_{jt}^{SP} \tag{30}$$

As mentioned before, $\Delta\epsilon_{jt}^{PC}$ and $\Delta\epsilon_{jt}^{SP}$ represent the exogenous changes in the wages (the exogenous productivity changes). Since, by definition, the Bartik instruments above are demand shifters, I can write:

$$\begin{cases} \Delta\epsilon_{jt}^{PC} & = \gamma_{pc}^{pc,inst} B_{jt}^{PC} + \gamma_{sp}^{pc,inst} B_{jt}^{SP} + \Delta\eta_{jt}^{PC} \\ \Delta\epsilon_{jt}^{SP} & = \gamma_{pc}^{sp,inst} B_{jt}^{PC} + \gamma_{sp}^{sp,inst} B_{jt}^{SP} + \Delta\eta_{jt}^{SP} \end{cases} \tag{31}$$

where $\Delta\eta_{jt}^{PC}$ and $\Delta\eta_{jt}^{SP}$ are unobserved changes in the productivity change which are, by construction, uncorrelated with the demand shocks captured by the Bartik instruments. I can then redefine the labor demand equations as:

$$\begin{cases} \Delta\text{comp}_{jt}^{PC} & = \gamma_{pc}^{pc} pc_{jt} + \gamma_{sp}^{pc} sp_{jt} + \gamma_{pc}^{pc,inst} \Delta B_{jt}^{PC} + \gamma_{sp}^{pc,inst} \Delta B_{jt}^{SP} + \Delta\eta_{jt}^{PC} \\ \Delta\text{comp}_{jt}^{SP} & = \gamma_{pc}^{sp} pc_{jt} + \gamma_{sp}^{sp} sp_{jt} + \gamma_{pc}^{sp,inst} \Delta B_{jt}^{PC} + \gamma_{sp}^{sp,inst} \Delta B_{jt}^{SP} + \Delta\eta_{jt}^{SP} \end{cases} \tag{32}$$

where, as before, the physician demand elasticities $\gamma_i^k \, \forall i, k = pc, sp$ are the parameters of interest. These are identified using changes in physician supply which are not correlated with $\Delta\eta_{jt}^{PC}$ and $\Delta\eta_{jt}^{SP}$. I use the interaction of the Bartik instruments (which are uncorrelated with $\Delta\eta_{jt}^{PC}$ and $\Delta\eta_{jt}^{SP}$ by construction) with region-based Medicare malpractice insurance reimbursement units. The malpractice reimbursement rates set by Medicare are set by procedure (MP) and location, based on observed malpractice insurance premia, and they therefore proxy for the total malpractice insurance cost paid by physicians. The location adjustments ($MPGPCI$), called geographical practice cost indices, which are set differently for procedure and malpractice insurance rates, react to the increase in the cost of living of the location. The use of changes in the malpractice insurance rates are useful as tiebreakers. To illustrate, take two cities with the same increase in physician demand, but that experience different changes in the total malpractice insurance costs. The city that experiences the higher cost increase will be less desirable to physicians, at an equal demand. This analysis will remain valid as long as the exclusion restrictions remain respected:

$$\mathbb{E}\left(\Delta\eta_{jt}^{PC} \Delta Z_{jt}\right) = 0 \tag{33}$$

$$\mathbb{E}\left(\Delta\eta_{jt}^{SP} \Delta Z_{jt}\right) = 0 \tag{34}$$

where $\Delta Z_{jt} \in \left\{\Delta B_{jt}^{k} \Delta MPGPCI_{jt},\, \Delta B_{jt}^{k} \Delta MP_{jt}^{k'} \, \forall k, k' = PC,\, SP\right\}$. Once again, recall that

$$\mathbb{E}\left(\Delta\eta_{jt}^{PC} \Delta B_{jt}^{PC}\right) = 0 \tag{35}$$

$$\mathbb{E}\left(\Delta\eta_{jt}^{SP} \Delta B_{jt}^{SP}\right) = 0 \tag{36}$$

by construction.

7 Results

7.1 Physician Supply

Table 4 presents the results for physicians' preferences in their choice of location, both for primary care physicians (first panel) and specialists (second panel). The first column shows the base results, while the top-50 and foreign columns report the differential results for top-50 and foreign-educated residents, respectively.

Table 4: Physician Supply: Individual Preferences ($\beta^{k,\ell}$, $\alpha^{k,\ell}$, $\beta^{k,\ell}_{state}$, $\beta^{k,\ell}_{HRR}$)

	Amenities		Amenities, Top 50		Amenities, Foreign	
	PC	SP	PC	SP	PC	SP
$\beta^{k,\ell}$	0.39	0.59	0.20	0.25	0.01	0.04
	(0.018)	(0.020)	(0.014)	(0.012)	(0.008)	(0.006)
	Income		Income, Top 50		Income, Foreign	
	PC	SP	PC	SP	PC	SP
$\alpha^{k,\ell}$	0.03	0.15	-0.005	-0.023	-0.001	0.001
	(1.18e-06)	(7.57e-06)	(2.58e-06)	(1e-05)	(0.001)	(0.001)
	State		State, Top 50			
	PC	SP	PC	SP		
$\beta^{k,\ell}_{state}$	2.77	1.75	0.42	-0.38		
	(0.012)	(0.035)	(0.018)	(0.053)		
	HRR		HRR, Top 50			
	PC	SP	PC	SP		
$\beta^{k,\ell}_{HRR}$	2.35	2.57	-1.48	-0.29		
	(0.005)	(0.042)	(0.008)	(0.063)		

Notes: These results come from the second specification of the physician supply analysis described in the paper. Magnitude of the α represents the elasticity of demand of a location with respect to income. Magnitude of the state and HRR coefficients represent the semielasticity of demand with respect to whether the choice is within the same state or area (HRR) of residency, respectively. The coefficients are obtained through maximum likelihood estimation of the conditional logit model based on individual-level data on residency and choice locations. The sample includes all residents finishing residency, by specialty, between 2012 and 2016. The choice set is given by the set of hospital referral regions in the United States as defined by Dartmouth Atlas.

In general, physicians like higher net incomes and higher amenities. Specialists are more elastic to income than primary care physicians. The fact that specialists appear to be more elastic to income is due to selection into specialty care before residency. Specialists are paid substantially more than primary care physicians in salary as well as for specialty procedures in fee-for-service settings. Since I do not directly model the choice of specialty group, which happens at the end of medical school with the choice of residency, individuals who are more interested in the monetary component would self-select into specialty care. Top-50 residents are less elastic to income, even though the estimate is quite small. Specialists are also more elastic to amenities than primary care physicians. Top-50 residents are more elastic to amenities. Recall that amenities also include investment in health facilities, and that cities tend to have both higher amenities and more research hospitals, where the top residents would often prefer to go, similarly to economics Ph.D.s preferring research universities. Finally, foreign-educated physicians are shown to not behave significantly differently from U.S.-educated physicians.

As mentioned beforehand, this estimation allows me to recover the unobserved amenities from the contraction mapping. The density distribution of the recovered unobserved amenities is shown in Figure 3. As expected, cities have higher unobserved amenities than rural places.

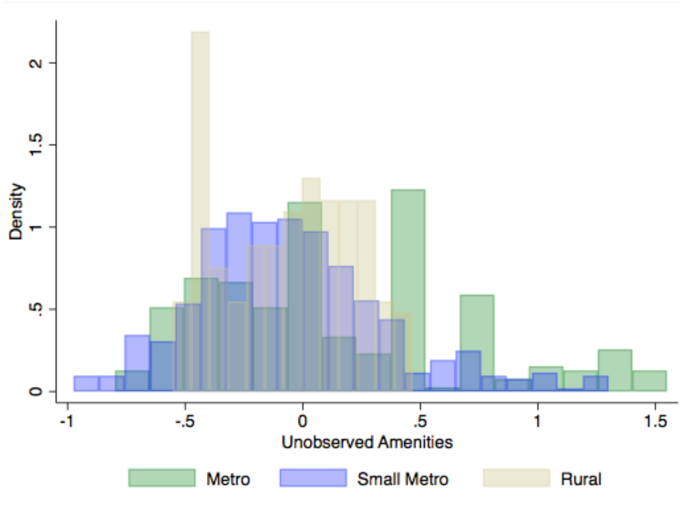

Figure 3: Density of Amenities, by Location Type

Notes: This figure shows the density distribution of the amenities implied by the model (both observed and unobserved), by location type. As expected, cities have higher amenity levels than rural areas.

Preference for retention

Table 4 also reports the results for physicians' preferences for retention in the last two rows. Preferences of physicians with the same demographic characteristics k, ℓ for a HRR j can differ due to preferences to remain within the residency's HRR and state. This is motivated by two facts in the health literature: first, many physicians that want to return to their state of birth tend to pick a residency that already fulfills their preference, so this preference also proxies for a preference to return to the birthplace; second, many physicians build personal and professional networks during residency that lead to their choice of remaining close to their residency location, for example because of a joint location decision with a spouse.[6] These two facts lead to a third, widely discussed fact: many states display high rates of residency retention, with 54 percent of all residents in the United States remaining within their state of residency for their first job. The dataset I create replicates this fact, with 51.4 percent of physicians picking to remain within the same state of residency.

Table 4 shows the estimates from the maximum likelihood estimation of the conditional logit model. The last two rows of Table 4 report the semi-elasticity of demand with respect to whether or not the choice is within the same state of residency and within the same hospital referral region, respectively. Generally, I find that all physicians exhibit strong preferences to remain close to their residency location. Top-50 primary care physician residents are more likely to be retained in the same state as residency, unlike specialty residents. Residents in all specialties are less likely to be retained in the same area as residency if they are from a top residency. I find that primary care physicians are about 3.8 times more likely to pick a job within the same state of residency and about 3.4 times more likely to pick a job within the same hospital referral region as the residency. By contrast, specialists are 2.8 times more likely to pick a job within the same state as residency and about 3.6 times more likely to pick a job within the same hospital referral region. I am able to reject that retention values can be the same between primary care physicians and specialists. SectionA.2.4 in the Appendix reports the semi-elasticity of retention estimates on a year-by-year basis.

[6]Note that, due to data availability, I am not able to control for spousal characteristics and/or to model spousal decisions.

I compare the base estimates with those of individuals that attended a top-50 residency. I find that top-50 residents in primary care are 0.4 times more likely to remain in the same state as residency, but they are 1.5 times less likely to remain in the same area as residency relative to the base estimates. By contrast, I find that top-50 residents in specialty care are 0.3-0.4 times less likely to be retained within the same state and area of residency relative to the base estimates.

This evidence for a preference for retention agrees with the previous survey literature on physicians' choice of location (including, but not limited to, Cooper 1975), supporting and strengthening the literature's survey findings.

Comparing these results to the labor literature, I find a very interesting result. While all physicians are clearly high-skilled workers, primary care physicians similar preference levels for retention as unskilled workers. Diamond (2015), for example, reports a base semi-elasticity of college of workers of being retained in their state of birth of about 2.6. That estimate is closer to the values I find for specialists, but much lower than the values that I find for primary care physicians. This shows that there are extremely important differences not only across occupation types, but also within occupations that might be ignored in current analyses.

Finally, I find extremely high estimates of retention of physicians that completed their residency rurally. Due to the small sample problem, the standard errors are more noisy, but I find that rural residents are at least one time more likely to be retained. This agrees with data evidence showing that, despite the fact that fewer than one percent of residencies are rural, physicians who complete a residency in a rural area are 70 percent more likely to pick a rural practice location.

7.2 Physician Demand

Table 6 presents the parameter estimates for physician demand. The estimated elasticity of labor substitution is equal to 1.01. The other specification allows for not only the labor substitution between the two types of physicians, but also for the effect of varying the employment of either type of physician on the productivity of either type of physician. I observe a negative own-elasticity of primary care physicians' wages, but a positive own-elasticity of specialists' wages, suggesting that the specialists "feed off each other." This finding agrees with the results of the labor literature for high-skilled individuals, as in Diamond (2016).

Finally, I observe a negative elasticity of specialists' compensation with respect to primary care employment, as suggested by the data evidence and discussion in Section 3.

8 Counterfactuals

Policymakers have historically used different monetary incentives to attract physicians to rural areas, most notably loan forgiveness and salary incentives. Note that the salary incentives are bonuses offered in addition to the already-higher equilibrium salaries. The average sign-up bonus for a physician to practice in rural areas has been estimated to be around $7,500.

I run three different counterfactuals under this setup: I evaluate the effect of current policies compared to no incentives, I evaluate the effect of a policy that switches the focus from loan forgiveness to salary incentives and offers the increased salary incentives to all physicians independently of their specialty, and I evaluate the effect of a policy that uses the current spending to increase salary incentives, but that offers such increased salary incentives to primary care physicians only.

Table 5: Physician Demand

	(1)	(2)
ρ	1.02	
	(0.004)	
γ_{pc}^{pc}		-0.23
		(0.06)
γ_{sp}^{pc}		0.24
		(0.05)
γ_{sp}^{sp}		0.19
		(0.05)
γ_{pc}^{sp}		-0.20
		(0.05)

Notes: These results come from the physician demand analysis described in the paper. ρ represents the elasticity of labor substitution between primary care physicians and specialists in specifications (1) and (2) in the model. All γs represent the reduced-form coefficient determining the relationship between the employment of either type of physicians and their total compensation.

First, I estimate that the combination of the current policies has led to 1.2 percent more physicians choosing rural areas. I estimate that loan forgiveness alone has increased the number of new primary care physicians choosing rural areas by 0.5 percent and of new specialists by 1.3 percent. Since loan forgiveness is particularly attractive to those students with high loans, a resorting effect happens, where residents re-maximize their utility and pick rural areas when their loans are the greatest, while those with lower loans move to the city as amenities offer a higher utility. Note that this is particularly important for those who has low or zero loans. Importantly, many physicians who hold no loans are the foreign-educated ones. About one-third of primary care physicians hold very low or no loans at all. The difference in the policy design is then key: loan forgiveness is equivalent to an incentive equal to the maximum of the loan forgiveness offered only if the physician in fact holds as high of an amount in loans as that amount. To illustrate, if a physician holds \$50,000 in loans and is offered \$100,000 in loan forgiveness, she can at most receive \$50,000 because that is the amount of student loans she holds. Salary incentives, by contrast, are the same for all. In other words, salary incentives target the full physician population, while loan forgiveness only attracts those with loans and is the more attractive, the higher the amount of loans held. Table 6 shows the percentage change in the primary care (PC) and specialty care (SP) physician population with respect to the base case of no incentives in any area.

Second, I add a \$7,500 bonus value in income by choosing rural areas. This differential effect, reported in the second row of Table 6, adds another 0.2 percent with respect to the primary care population and an another 0.1 percent with respect to the specialist population in the world with only loan forgiveness.

These two exercises have replicated the effect of currently implemented policies. Regardless of specialty, the two current policies have led to a 1.2 percent increase in the number of physicians picking rural areas. Results suggest that specialists respond the most to current policies, which is mainly driven by the focus on loan forgiveness and the very high loans usually held by specialists. Note that a specialist identical in every other way to a primary care physician will hold more loans than that primary care physician even if they attended the same medical school. This fact is simply due to the length of residency, which is longer for specialists than for primary care physicians, and during which interest keeps accruing and leads to a higher

Table 6: Counterfactual Effect, by Specialty Group

	PC	SP
Loan Forgiveness	0.7 percent	1.3 percent
Salary Incentives	0.2 percent	0.1 percent

Notes: The table reports the percentage changes in physician population, by specialty type. "Loan Forgiveness" only includes the effect of a loan forgiveness policy, "Salary Incentives" includes the additional effect of a $7,500 salary incentives.

initial loan amount for specialists to repay.

The detailed information fed into my model also allows me to analyze the average quality level of the physicians who are incentivized by these policies. I find that while loan forgiveness might appear as an attractive option to bring physicians to rural areas, the physicians attracted there are in the bottom 25 percent of the quality ladder. This characteristic is once again due to the fact that the physicians who are most responsive to loan forgiveness are those with very high debt. There are two types of physicians in the United States who are more likely to hold the most loans: those who attended top private medical schools and those who attended private low-ranked medical schools. While attending top medical schools often leads to high student loans, it also gives physicians access to top residencies and, subsequently, to the most remunerative jobs. By contrast, private low-ranked medical schools also lead to high student debt, but they are often followed by low-ranked residencies and low-quality jobs. Physicians who face this situation then find themselves with high debt and low-quality job perspectives, which incentivizes them to accept loan forgiveness whenever possible to reduce their student loan burden.

Given this, I run the experiment of using the U.S. spending currently allocated to medical-student-loan forgiveness, which is currently between $350 and $400 million, to increase the salary incentives offered for rural area employment across all physicians. The evaluation of current and alternative policies is presented in Table 7. While this experiment maintains the same level of spending, salary incentives incentivize all physicians equally, not just those with high loans. The redistribution of this loan-forgiveness spending to rural salary incentives across all physicians leads to an extra bonus of roughly $35,000 for each physician picking rural areas. I find that such a policy would lead to almost six times more primary care physicians picking rural areas than the number of physicians currently picking rural areas because of current policies (i.e., the current policy effect described above). As discussed above, the fact that primary care physicians respond to this alternative policy a lot more is due to the loan distribution across the physician population: the switch to higher salary incentives now offers an incentive for all, including those physicians who hold low-to-no loans, and this incentive now simply leads to a trade-off between higher salary or higher amenities, where the higher salary wins for quite a few more primary care physicians. Since primary care physicians are the main physician category that currently provides medical care to rural areas, and since they do not need a particular technological infrastructure to do so, these results suggest that policymakers should retarget spending from loan forgiveness to salary incentives. The switch to salary incentives also allows for a more varied pool of physicians in terms of quality ranking, increasing the overall average quality of physicians.

Finally, considering that primary care physicians seem to respond well to increased salary incentives, offering such incentives to primary care physicians only might be even more effective. Therefore, I run the same exercise of switching to salary incentives, but targeting primary care physicians only. The increase in the primary care physician population is now substantial, while a few specialists are lost compared to the "no

Table 7: Impact of Current and Alternative Policies, by Specialty Group

Policy Environment	PC	SP
Data: New Rural Physicians	1,039	1,060
Model		
Current Policies	1,039	1,060
Remove All Incentives	971	773
Effect of Current Policies	+68	+287
Effect of Alternative Policy 1 - Target All	+407	+132
Effect of Alternative Policy 2 - Target Primary Care Only	+1,029	-22

Notes: The table reports the impact of current and alternative policies on the physician population, by specialty type.

incentives" world, due to the increased competition coming from the rising primary care population. Note that in reality, it is difficult to believe that hospitals would not rationally respond to such a high increase in incentives (now approaching $100,000) by adjusting their wage offers down. Since hospitals are not modeled in a very advanced way in this final counterfactual scenario, this second result should be taken with a pinch of salt. This is not true for the first alternative scenario, where instead the value of the incentives for any given physician remains very similar, but is simply redistributed across the population.

These exercises served as an example to show how monetary incentives are currently not strong enough for physicians to move to rural areas. They also suggest that policies aimed at the retargeting of spending from loan forgiveness to salary incentives would lead to almost 6 times more primary care physicians choosing rural areas. The pool of physicians attracted would be of a better quality on average under salary incentives compared to current policies because the full population of physicians, independently of the amount of loans held, would face the same incentives. Finally, the number of physicians attracted is still not solving the shortage very quickly. The reason for this is physicians' preference to remain close to their residency location, which makes their location-switching costs extremely high. Therefore, a point that is indirectly unveiled by this analysis is that investing in rural residencies might be the most effective strategy to bring physicians to rural areas. This point is also supported by the fact that physicians who completed residency in a rural area are at least one time more likely to pick a practice location in a rural area.

9 Conclusion

This paper used a combination of detailed and novel data to be able to assess the main factors affecting physicians' choice of location of their first job following residency.

The results suggest that physicians respond positively to compensation and amenities, as theory would suggest. However, physicians do not respond much to monetary incentives. Specialists are more elastic than primary care physicians and respond more to both higher compensation and higher amenities. Nevertheless, none of these characteristics matter as much as the location of their residency. The retention effect is strong both in the data, which show that those retained all the way from medical school have less than a one-in-three chance of moving away after residency, and in the retention estimates, which show that all physicians prefer remaining close to their residency location.I find that primary care physicians are about 3.5 times more likely to pick a job within the same state of residency and 4 times more likely to pick a job within the same hospital referral region as the residency. On the other hand, specialists are about 2.5 times more likely to pick a job within the same state as residency and about 3 times more likely to pick a job within the same hospital referral region. I also find that top-50 residents are more difficult to be retained. The differences in preferences between primary care physicians and specialists, including their differences in how much they value remaining close to their residency location, illustrates that there are key within-occupation differences that are often ignored but that have important implications for policy design.

The results suggest that short-term policies should incorporate the fact that physicians respond to factors other than compensation and that amenities and residency location play key roles in their choices. I use the model to analyze the performance of current policies targeted at bringing physicians to rural areas and find that the combination of current policies has led to 1.2 percent more physicians in rural areas. In particular, I find that 0.5 percent more primary care physicians and 1.3 percent more specialists have picked rural areas due to loan forgiveness alone. Monetary incentives in the form of salary payments averaging $7,500 are responsible for a further 0.2 percent increase in primary care physicians and 0.1 percent increase in specialists. By retargeting the spending currently used for loan forgiveness to increase salary incentives for rural employment, I find that almost six times more primary care physicians would pick rural areas than the number of primary care physicians currently picking rural areas because of the current policies. Since primary care physicians are the main physician category that currently provides medical care to rural areas, and since they do not need a particular infrastructure to do so, these results suggest that policymakers should retarget spending from loan forgiveness to salary incentives and that offering salary incentives to primary care physicians only could be more effective. The average quality of the physicians attracted under these higher salary incentives is also better compared to the quality distribution under the current policies.

Despite these results, the number of physicians attracted to rural areas remains small even under these alternative policies. The reason for this result lies in the strong preference that physicians have to remain close to their residency location. Therefore, the mechanism that indirectly appears to be effective at bringing physicians to rural places seems to be the following: residencies need to be made available, offered, and maintained rurally, since currently fewer than one percent of all residencies are offered in rural areas. Not only it is much easier to retain than to attract physicians to a new location, but physicians that complete their residencies in rural places are at least one time more likely to locate in a rural area. Since the residency choice is not directly modeled in this paper, these facts are only indirect results from my analysis, but will be addressed in future work.

References

1. Acemoglu, Daron, Amy Finkelstein, and Matthew J. Notowidigdo. "Income and Health Spending: Evidence from Oil Price Shocks." *Review of Economics and Statistics* 95.4 (2013): 1079-1095.
2. Agarwal, Nikhil. "An Empirical Model of the Medical Match." *American Economic Review* 105.7 (2015): 1939-78.
3. American Hospital Association. "Physician Ownership and Self-Referral in Hospitals: Research on Negative Effects Grows." *Trendwatch.* (April 2008).
4. Anderson, Gerard F., et al. "Health Spending in the United States and the Rest of the Industrialized World." *Health Affairs* 24.4 (2005): 903-914.
5. Armour, Brian S., et al. "The Effect of Explicit Financial Incentives on Physician Behavior." *Archives of Internal Medicine* 161.10 (2001): 1261-1266.
6. Autor, David H., and David Dorn. "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market." *American Economic Review* 103.5: 1553-97.
7. Baker, Laurence C. "The Effect of HMOs on Fee-For-Service Health Care Expenditures: Evidence from Medicare." *Journal of Health Economics* 16.4 (1997): 453-481.
8. Barro, Jason, and Nancy Beaulieu. "Selection and Improvement: Physician Responses to Financial Incentives." *National Bureau of Economic Research Working Paper* 10017 (2003).
9. Berry, Christopher R., and Edward L. Glaeser. "The Divergence of Human Capital Levels Across Cities." *Papers in Regional Science* 84.3: 407-44.
10. Berry, Steven, James Levinsohn, and Ariel Pakes. "Automobile Prices in Market Equilibrium." *Econometrica: Journal of the Econometric Society* (1995): 841-890.
11. Blomqvist, Ake, and Colin Busby. *How to Pay Family Physicians: Why 'Pay Per Patient' Is Better than Fee for Service.* C.D. Howe Institute Commentary 365 (2012).
12. Broomberg, J., and M. R. Price. "The Impact of the Fee-For-Service Reimbursement System on the Utilisation of Health Services. Part I. A Review of the Determinants of Physicians' Practice Patterns." *South African Medical Journal* 78.3 (1990): 130-132.
13. Card, David. "Immigration and Inequality." *American Economic Review* 99.2: 1-21.
14. Card, David, and Thomas Lemieux. "Can Falling Supply Explain the Rising Return to College for Younger Men? A Cohort-Based Analysis." *Quarterly Journal of Economics* 116.2: 705-46.
15. Chaix-Couturier, Carine, et al. "Effects of Financial Incentives on Medical Practice: Results from a Systematic Review of the Literature and Methodological Issues." *International Journal for Quality in Health Care* 12.2 (2000): 133-142.
16. Clemens, Jeffrey, and Joshua D. Gottlieb: "Do Physicians' Financial Incentives Affect Medical Treatment and Patient Health?." *The American Economic Review* 104.4 (2014): 1320-1349.
17. —. "In the Shadow of a Giant: Medicare's Influence on Private Physician Payments." *Journal of Political Economy* 125.1 (2017): 1-39.
18. Colas, Mark. "Dynamic Responses to Immigration." *Opportunity & Inclusive Growth Insitute Working Paper* 6 (2018).
19. Cooper, James K., et al. "Rural or Urban Practice: Factors Influencing the Location Decision of Primary Care Physicians." *Inquiry* 12.1 (1975): 18-25.

20. Cooper, Richard A., et al. “Economic and Demographic Trends Signal an Impending Physician Shortage.” *Health Affairs* 21.1 (2002): 140-154.

21. Croxson, Bronwyn, Carol Propper, and Andy Perkins. “Do Physicians Respond to Financial Incentives? UK Family Physicians and the GP Fundholder Scheme.” *Journal of Public Economics* 79.2 (2001): 375-398.

22. Diamond, Rebecca. “The Determinants and Welfare Implications of US Workers’ Diverging Location Choices by Skill: 1980-2000.” *American Economic Review* 106.3 (2016): 479-524.

23. Dunne, Timothy, et al. “Entry, Exit, and the Determinants of Market Structure.” *RAND Journal of Economics* 44.3 (2013): 462-487.

24. Emanuel, Ezekiel J., and Victor R. Fuchs. “The Perfect Storm of Overutilization.” *Journal of the American Medical Association* 299.23 (2008): 2789-2791.

25. Falcettoni, Elena. *The Consequences of Medicare Pricing: An Explanation of Treatment Choice*. Working Paper (2018).

26. Finkelstein, Amy. “The Aggregate Effects of Health Insurance: Evidence from the Introduction of Medicare.” *Quarterly Journal of Economics* 122.1 (2007): 1-37.

27. Fisher, Elliott S., et al. “The Implications of Regional Variations in Medicare Spending. Part 1: The Content, Quality, and Accessibility of Care.” *Annals of Internal Medicine* 138.4 (2003): 273-287.

28. Gawande, Atul, and BBC Audiobooks America. *Better: A Surgeon’s Notes on Performance*. Picador (Jan. 22, 2008).

29. Ginsburg, Paul B. “Fee-For-Service Will Remain a Feature of Major Payment Reforms, Requiring More Changes in Medicare Physician Payment.” *Health Affairs* 31.9 (2012): 1977-1983.

30. Gosden, Toby, et al. *Capitation, Salary, Fee-For-Service and Mixed Systems of Payment: Effects on the Behaviour of Primary Care Physicians*. The Cochrane Library (2000).

31. Grant, Darren. “Physician Financial Incentives and Cesarean Delivery: New Conclusions from the Healthcare Cost and Utilization Project.” *Journal of Health Economics* 28.1 (2009): 244-250.

32. Gruber, Jonathan, and Maria Owings. ”Physician Financial Incentives and Cesarean Section Delivery.” *National Bureau of Economic Research Working Paper* 4933 (1994).

33. Grumbach, Kevin, et al. “Primary Care Physicians’ Experience of Financial Incentives in Managed-Care Systems.” *New England Journal of Medicine* 339.21 (1998): 1516-1521.

34. Hall, Robert E., and Charles I. Jones. “The Value of Life and the Rise in Health Spending.” *Quarterly Journal of Economics* 122.1 (2007): 39-72.

35. Hellinger, Fred J. “The Impact of Financial Incentives on Physician Behavior in Managed Care Plans: A Review of the Evidence.” *Medical Care Research and Review* 53.3 (1996): 294-314.

36. Hemenway, David, et al. “Physicians’ Responses to Financial Incentives: Evidence from a For-Profit Ambulatory Care Center.” *New England Journal of Medicine* 322.15 (1990): 1059-1063.

37. Hendee, William R., et al. “Addressing Overutilization in Medical Imaging.” *Radiology* 257.1 (2010): 240-245.

38. Hickson, Gerald B., William A. Altemeier, and James M. Perrin. “Physician Reimbursement by Salary or Fee-For-Service: Effect on Physician Practice Behavior in a Randomized Prospective Study.” *Pediatrics* 80.3 (1987): 344-350.

39. Hillman, Alan L. “Financial Incentives for Physicians in HMOs.” *New England Journal of Medicine* 317.27 (1987): 1743-1748. APA

40. Hillman, Alan L., Mark V. Pauly, and Joseph J. Kerstein. “How Do Financial Incentives Affect Physicians’ Clinical Decisions and the Financial Performance of Health Maintenance Organizations?.” *New England Journal of Medicine* 321.2 (1989): 86-92.

41. Jacobson, Mireille, et al. “Does Reimbursement Influence Chemotherapy Treatment for Cancer Patients?.” *Health Affairs* 25.2 (2006): 437-443.

42. Katz, Lawrence F., and Kevin M. Murphy. 1992. “Changes in Relative Wages, 1963-1987: Supply and Demand Factors.” *Quarterly Journal of Economics* 107.1: 35-78.

43. Kirch, Darrell G., Mackenzie K. Henderson, and Michael J. Dill. “Physician Workforce Projections in an Era of Health Care Reform.” *Annual Review of Medicine* 63 (2012): 435-445.

44. Kulka, Amrita, and Dennis McWeeny. “Rural Physician Shortages and Policy Intervention.” (2018).

45. Lee, Sanghoon. “Ability Sorting and Consumer City.” *Journal of Urban Economics* 68.1 (2010): 20-33.

46. Leonardson, Gary, Rene Lapierre, and Doneen Hollingsworth. “Factors Predictive of Physician Location.” *Journal of Medical Education* (1985).

47. Levin, David C., and Vijay M. Rao. “Turf Wars in Radiology: The Overutilization of Imaging Resulting from Self-Referral.” *Journal of the American College of Radiology* 5.7 (2008): 806-810.

48. McElduff, P., et al. “Will Changes in Primary Care Improve Health Outcomes? Modelling the Impact of Financial Incentives Introduced to Improve Quality of Care in the UK.” *Quality and Safety in Health Care* 13.3 (2004): 191-197.

49. Melichar, Lori. “The Effect of Reimbursement on Medical Decision Making: Do Physicians Alter Treatment in Response to a Managed Care Incentive?.” *Journal of Health Economics* 28.4 (2009): 902-907.

50. Moretti, Enrico. “Real Wage Inequality.” *American Economic Journal: Applied Economics* 5.1: 65-103.

51. Nevo, Aviv. “Empirical Models of Consumer Behavior.” *Annual Review of Economics* 3.1: 51-75.

52. —. “Measuring market power in the ready-to-eat cereal industry.” *Econometrica* 69.2 (2001): 307-342.

53. —. “A Practitioner’s Guide to Estimation of Random-Coefficients Logit Models of Demand.” *Journal of Economics & Management Strategy* 9.4 (2000): 513-548.

54. Pham, Hoangmai H., and Paul B. Ginsburg. “Unhealthy Trends: The Future of Physician Services.” *Health Affairs* 26.6 (2007): 1586-1598.

55. Rasmusen, Eric. *The BLP Method of Demand Curve Estimation in Industrial Organization*. Department of Business Economics and Public Policy, Kelley School of Business, Indiana University. (2007).

56. Rosenblatt, Roger A., and L. Gary Hart. “Physicians and Rural America.” *Western Journal of Medicine* 173.5 (2000): 348.

57. Safran, Dana Gelb, Alvin R. Tarlov, and William H. Rogers. “Primary Care Performance in Fee-For-Service and Prepaid Health Care Systems: Results from the Medical Outcomes Study.” *Journal of the American Medical Association* 271.20 (1994): 1579-1586.

58. Scott, Anthony, et al. “The Effect of Financial Incentives on the Quality of Health Care Provided by Primary Care Physicians.” *Cochrane Database of Systematic Reviews* 20.1.9 (2010).

59. Shrank, William, et al. “Effect of Physician Reimbursement Methodology on the Rate and Cost of Cataract Surgery.” *Archives of Ophthalmology* 123.12 (2005): 1733-1738.

60. Stearns, Sally C., Barbara L. Wolfe, and David A. Kindig. “Physician Responses to Fee-For-Service and Capitation Payment.” *Inquiry* (1992): 416-425.

61. Steele, Henry B., and Gaston V. Rimlinger. “Income Opportunities and Physician Location Trends in the United States.” *Economic Inquiry* 3.2 (1965): 182-194.

62. United States Census Bureau. Census 2010 PMSA Definition:

https://www.census.gov/population/estimates/metro-city/99mfips.txt

63. Vincent, David W. “The Berry–Levinsohn–Pakes Estimator of the Random-Coefficients Logit Demand Model.” *Stata Journal* 15.3 (2015): 854-880.

Appendix

A.1 Physicians in Different Locations Act Differently

This subsection briefly summarizes some data findings in Falcettoni (2018), which support the hypothesis held in this paper that physicians take into consideration the level of reimbursement in their decision-making. Since physicians perform different procedures in different places, the choice of location inherently encompasses that information.

In particular, primary care physicians are able to perform more remunerative specialty procedures in rural areas, where specialists are not as present. Since specialty procedures have been paid increasingly more over the years, the possibility of carrying such procedures out in more rural areas makes rural areas more attractive from a remuneration point of you. However, classical analysis only considers a tradeoff between wages and amenities, and does not take into consideration this other very important remuneration channel.

I report here some data evidence on physicians acting differently along the urbanity index. In first need to identify what constitutes a specialty procedure. To do so, I look at the number of services per procedure carried out by each physician and see how many physicians in primary care perform it and how many specialists perform it over the entire dataset. Then, for each procedure, I calculate the percentage of services performed by primary care physicians versus specialists. I then consider the procedures of interest to be those performed by specialists 50-80 percent of the time and by the primary care the remaining 20-50 percent. Robustness checks show that the results are robust independently of the range chosen, but their effect is stronger for tighter ranges. This approach creates a specialization index for each procedure, from 0 to 1. An index value equal to 0 means that the procedure is only carried out by specialists, while an index value of 1 means that it is always carried out by primary care physicians. Therefore, the lower the index value, the more the procedure is a specialized one. Next, I consider all the procedures carried out by each physician and their respective specialization index values. I then take the average of these values across all the procedures carried out by each physician, for every physician. This generates a physician-level specialization index which marks whether or not each physician behaves as a specialist. Similarly to before, if a physician had an average of 0, it would mean that he only carried out specialty procedures and if he had an average of 1, it would mean that he only carried out PC procedures. Therefore, the lower the average, the higher the number of specialty procedures carried out by the physician. I call this variable the degree of specialization of physicians. I then focus on procedures carried out by PC physicians 20-50 percent of the time and look at the distribution of physicians across the urbanity index, as shown in Figure A.1.

I then restrict my attention to highly specialized procedures, i.e. the procedures carried out by specialists 70-80 percent of the time. I then analyze how many physicians and how many of these services are carried out by primary care physicians and by specialists across the urbanity level, and present the results in Figure A.2.

This brief data analysis shows how primary care physicians take over more specialty procedures in rural places, and do so even more in locations where specialists are fewer. The income differential between procedures typically carried out by primary care physicians and those typically carried out by specialists is key in this behavior and supports the hypothesis of this paper that reimbursements should be included when analyzing physicians' location choice.

For a more complete analysis on physicians' response to the financial incentives generated by this behavior, refer to Falcettoni (2018).

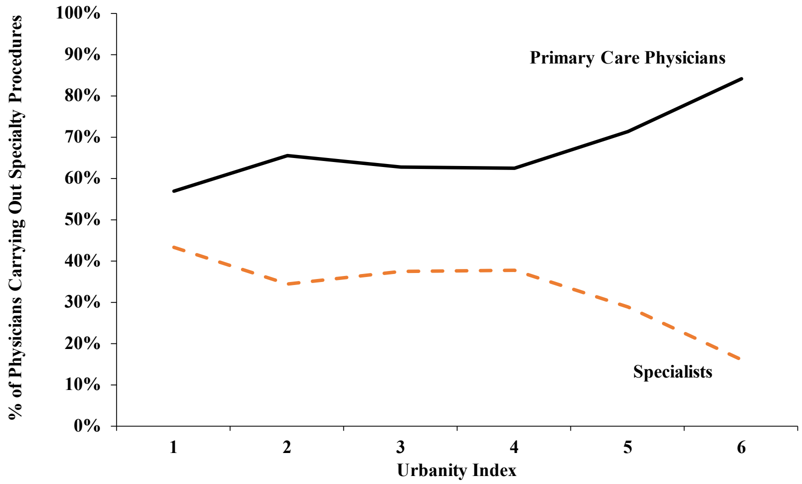

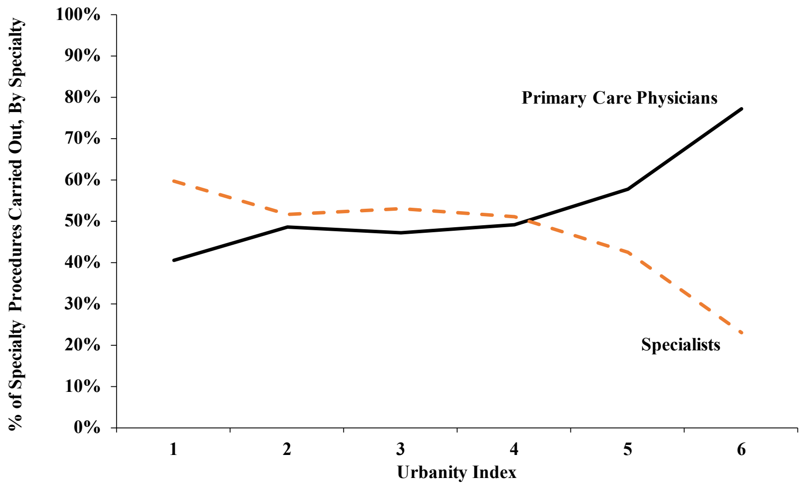

Figure A.1: Procedures carried out by specialists 50-80% of the time

Notes: This figure concentrates on procedures carried out by specialists 50-80% of the time. The first figure looks at the percentage of physicians carrying out these procedures (even just once) who are in primary care vs some specialty, across urbanity levels. The second figure looks at the percentage of services provided by primary care vs some specialty, across urbanity levels. The urbanity of the area is an index from 1 to 6, where a higher value denotes a more rural area. Source: Author's calculations based on data from CMS.

A.2 Net income in the mean utility $\delta_j^{k,\ell}$

A.2.1 Model

In this section, I calculate the average income for a location by specialty and demographics to incorporate compensation in the mean utility. Therefore, physician i in specialty k and with demographic characteristics of type ℓ then solves:

$$\max_j u_{ij} \quad = \quad \overbrace{\alpha^{k,\ell} y_j + \beta^{k,\ell} x_j + \xi_j^{k,\ell}}^{\delta_j^{k,\ell}} + \overbrace{\beta_j^{k,\ell} x_{ij}}^{\mu_{ij}} + \epsilon_{ij} \tag{37}$$

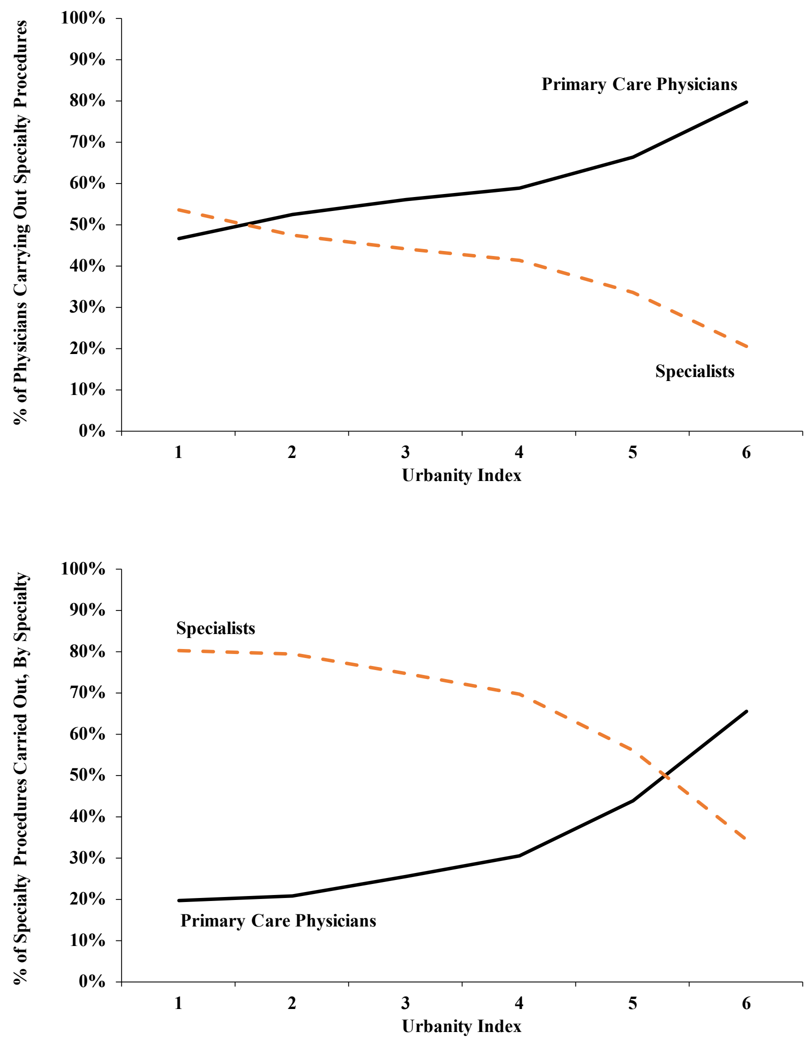

Figure A.2: Procedures carried out by specialists 70-80% of the time

Notes: This figure concentrates on procedures carried out by specialists 70-80% of the time. The first figure looks at the percentage of physicians carrying out these procedures (even just once) who are in primary care vs some specialty, across urbanity levels. The second figure looks at the percentage of services provided by primary care vs some specialty, across urbanity levels. The urbanity of the area is an index from 1 to 6, where a higher value denotes a more rural area. Source: Author's calculations based on data from CMS.

where (y_j, x_j) are now the location characteristics (net income as defined in Section 3.1 and amenities, respectively), while everything else follows exactly as above. Bearing in mind that now $\delta_j^{k,\ell}$ contains compensation as well, everything else follows the same setup as in the main text.

A.2.2 Estimation

Here, I briefly show what varies in the estimation procedure if I allow for an average compensation by location that is invariant at the individual level, so that each type observes the same compensation offer.

Let $\mathbb{I}_{ij}$ be a dummy variable that takes value one if physician i chooses location j. Recall that the probability that a physician i of type ℓ in specialty k picks location j in a given year is:

$$\hat{s}_{ijt}^{k,\,\ell} = \frac{\exp\left\{\delta_{jt}^{k,\,\ell} + \beta_j^{k,\,\ell} x_{ijt} + \epsilon_{ijt}\right\}}{\sum_{m=1}^{M} \exp\left\{\delta_{mt}^{k,\,\ell} + \beta_m^{k,\,\ell} x_{imt} + \epsilon_{imt}\right\}} \tag{38}$$

The discussion above on the log-likelihood remains valid here, therefore I use the contraction mapping from BLP (1995) as above, so that for every $\boldsymbol{\beta_j^{k,\,\ell}}$, there exists a unique $\boldsymbol{\delta_j^{k,\,\ell}} = \left(\delta_1^{k,\,\ell}, ..., \delta_J^{k,\,\ell}\right)$ that matches the observed and estimated shares. The log-likelihood is then given by:

$$\mathcal{L}^{k,\,\ell}\left(\boldsymbol{\delta^{k,\,\ell}}\left(\boldsymbol{\beta_j^{k,\,\ell}}\right), \boldsymbol{\beta_j^{k,\,\ell}}\right) = \frac{1}{N_\ell^k} \sum_{i=1}^{N_\ell^k} \sum_{j=1}^{J} \mathbb{I}_{ij} \log\left(\hat{s}_{ijt}^{k,\,\ell}\left(\boldsymbol{\delta^{k,\,\ell}}\left(\boldsymbol{\beta_j^{k,\,\ell}}\right), \boldsymbol{\beta_j^{k,\,\ell}};\, x_{ij}\right)\right) \tag{39}$$

Now that both $\beta_j^{k,\,\ell}$ and $\boldsymbol{\delta_j^{k,\,\ell}} = \left(\delta_1^{k,\,\ell}, ..., \delta_J^{k,\,\ell}\right)$ are recovered, all that is left to be estimated are the parameters on the net income $\alpha^{k,\,\ell}$ and on the amenities $\beta^{k,\,\ell}$.

The key issue in this part of the estimation is the endogeneity caused by the unobserved amenities ξ_j. While I include many amenities in the amenity index I discussed in Section 6.1, there could be other unobserved amenities that still impact physicians' choices. One such example is the fact that physicians might pick places surrounded by people that are similar to them. Due to the fact that physicians are paid more where amenities are fewer and that procedure revenue varies geographically as summarized in Section A.1 above, the correlation between net income and unobserved amenities is strictly negative, i.e.

$$\mathbb{E}\left[y_j \xi_j^{k,\,\ell}\right] < 0 \tag{40}$$

This would bias $\alpha^{k,\,\ell}$ downward, as it can easily be verified by running a simple OLS regression, which estimates negative coefficients on income.

I utilize the income instrument discussed in Section 6.2 to address this issue. The identifying assumption is that there is enough variation in the procedure mix carried out in different places and that these differences in the procedure mix are uncorrelated with location unobservables. In formal notation:

$$\mathbb{E}\left[B_{jt}^k \xi_j^{k,\,\ell}\right] = 0\, \forall k = PC,\, SP \tag{41}$$

A.2.3 Results

As mentioned earlier in the text, I also run the model again using the micro data to calculate the average income that a physician would receive in a given location to compare it to the existing literature. I then run two different specifications of the model in which net income is a component of the mean utility. The two specifications, similarly to the main text, only vary along the demand component. First, I assume that the demand elasticities are only a function of the elasticity of labor substitution between primary care physicians and specialists (Model 1). In the second specification, I relax this assumption and I let compensation respond to the employment of either type of physician (Model 2). The results are shown to be qualitatively robust, indicating that a comprehensive measure of income is the main driver for the results.

Table A.1 presents the results for physicians' preferences in their choice of location, both for primary care (PC) and specialists (SP). The first column for both models reports the base results. The second and third column report the differential estimates for top-50 and foreign-educated residents. In general, physicians like higher net incomes and higher amenities. Top-50 residents are more elastic to both, while I do not observe significant differences between U.S.- and foreign-educated physicians. Specialists are more elastic than primary care physicians with respect to both factors. As mentioned beforehand, the two specifications in Table A.1 only vary on the demand side. I find that the elasticity that primary care physicians display is very similar in both cases. Specialists' elasticity to income is about double the elasticity found when the coefficient is estimated using the individual-level data. I find that primary care physicians are about 3.7 times more likely to pick a job within the same state of residency and about 3.4 times more likely to pick a job within the same hospital referral region as the residency. By contrast, specialists are 3 times more likely to pick a job within the same state as residency and about 2.8 times more likely to pick a job within the same hospital referral region. I am able to reject that retention values can be the same between primary care physicians and specialists. The estimates of retention values indicate that top-50 residents are less likely to be retained in the same state or area of residency. The next subsection reports the semi-elasticity of retention estimates on a year-by-year basis.

Table A.2 reports the estimates for the demand side. I find qualitatively-similar results for both specifications of the model relative to the estimates reported in the main text.

A.2.4 Year-by-Year Preference for Retention

Table A.3 reports the results for physicians' preferences for retention where I let the estimate vary on a year-by-year basis. The retention estimates are very robust in this case as well, despite being more noisy because of the smaller sample available on a year-by-year basis.

The estimated retention value of staying within the same hospital referral region has increased 11 percent between 2012 and 2016, suggesting that primary care physicians value retention more and more. The value of staying within the same state for primary care physicians has also displayed an increase over the panel period analyzed, by a little over 19 percent. Specialists' preference to stay within the same state as residency has remained fairly constant across the panel.

I compare the base estimates with those of individuals that attended a top-50 residency. These last estimates seem to be more noisy, due to the small selected sample. I cannot conclude that top-50 primary care residents have a different value to remain within the same state of residency, but top-50 primary care residents have a lower value to remain within the same hospital referral region. Top-50 specialty residents display a lower value of retention, and they are about 1-1.5 times less likely to be retained at the hospital referral region- and state-level, across all years.

Table A.1: Physician Supply: Income in the Mean Utility ($\delta_{j}^{k,\ell}{}_{t}$)

	(1 - Base)		(1 - Top 50)		(1 - Foreign)	
	PC	SP	PC	SP	PC	SP
Compensation	0.043	0.074	0.019	0.020	-0.001	0.001
	(0.002)	(0.003)	(0.001)	(0.001)	(0.001)	(0.001)
Amenities	0.39	0.59	0.20	0.25	0.01	0.04
	(0.018)	(0.020)	(0.014)	(0.012)	(0.008)	(0.006)
Hansen's J stat	172.23					
p-value	0.3339					
	(2 - Base)		(2 - Top 50)		(2 - Foreign)	
	PC	SP	PC	SP	PC	SP
Compensation	0.036	0.065	0.016	0.015	-0.0004	-0.001
	(0.002)	(0.003)	(0.002)	(0.001)	(0.0005)	(0.001)
Amenities	0.49	0.73	0.24	0.29	0.001	0.047
	(0.019)	(0.026)	(0.014)	(0.013)	(0.006)	(0.006)
$\beta^{k,\ell}_{state}$	2.71	2.01	-0.35	-0.23		
	(0.041)	(0.037)	(0.069)	(0.050)		
$\beta^{k,\ell}_{HRR}$	2.40	1.79	-0.08	-0.36		
	(0.046)	(0.041)	(0.079)	(0.058)		
Hansen's J stat	143.95					
p-value	0.462					

Notes: These results come from the physician supply analysis described in the paper. Magnitude of all rows but the last two represents the elasticity of the mean utility to each variable, by specialty group. Magnitude of the last two rows represents the semielasticity of demand with respect to whether the choice is within the same state or area (HRR) of residency, respectively. The coefficients on individual preferences and the mean utility levels are obtained through maximum likelihood estimation of the conditional logit model based on individual-level data on residency and choice locations. The coefficients on the mean utility are obtained through two-step generalized method of moments. The coefficients for top-50 and foreign residents are the differential effects of residents that graduated from a top-50 place and from being foreign with respect to the base coefficients. The sample includes all residents finishing residency, by specialty, between 2012 and 2016. The choice set is given by the set of hospital referral regions in the United States as defined by Dartmouth Atlas. Compensation is net income defined as revenues coming from salary and reimbursements minus the expenses coming from malpractice insurance, rent, and student loan repayment. Health professional shortage areas that offer loan forgiveness exclude student loan repayment. I also let the individual preference parameters vary by year in a robustness check. The results from this estimation are available in the appendix.

Table A.2: Physician Demand

	(1)	(2)
ρ	1.01	
	(0.004)	
γ_{pc}^{pc}		-0.55
		(0.06)
γ_{sp}^{pc}		0.51
		(0.05)
γ_{sp}^{sp}		0.16
		(0.05)
γ_{pc}^{sp}		-0.19
		(0.05)

Notes: These results come from the physician demand analysis described in the paper. ρ represents the elasticity of labor substitution between primary care physicians and specialists in specifications (1) and (2) in the model. All γs represent the reduced-form coefficient determining the relationship between the employment of either type of physicians and their total compensation. All coefficients from the physician demand analysis are jointly estimated with the physician supply parameters of the main utility through two-step GMM.

Table A.3: Physician Supply: Preferences for Retention ($\beta_{state}^{k,\ell}$, $\beta_{HRR}^{k,\ell}$)

	State		State, Top 50		HRR		HRR, Top 50	
	PC	SP	PC	SP	PC	SP	PC	SP
2012	2.17	1.84	0.65	-0.32	2.79	2.11	-1.38	-0.26
	(0.08)	(0.06)	(0.41)	(0.12)	(0.09)	(0.06)	(0.50)	(0.13)
2013	2.14	1.56	0.52	-0.53	2.83	2.02	0.13	-0.14
	(0.08)	(0.06)	(0.42)	(0.15)	(0.09)	(0.07)	(0.70)	(0.16)
2014	2.54	1.60	-1.02	-0.59	2.81	2.10	0.58	-0.23
	(0.07)	(0.06)	(0.35)	(0.14)	(0.08)	(0.07)	(0.45)	(0.15)
2015	2.63	1.75	-0.33	-0.31	2.78	2.01	-0.67	-0.60
	(0.08)	(0.06)	(0.29)	(0.14)	(0.09)	(0.07)	(0.40)	(0.16)
2016	2.59	1.79	0.72	-0.45	3.10	2.10	-0.95	-0.43
	(0.08)	(0.06)	(0.42)	(0.13)	(0.09)	(0.07)	(0.45)	(0.15)

Notes: These results come from the physician supply analysis described in the paper. Magnitude represents the semielasticity of demand with respect to whether the choice is within the same state or area (HRR) of residency, respectively. The coefficients are obtained through maximum likelihood estimation of the conditional logit model based on individual-level data on residency and choice locations. The sample includes all residents finishing residency, by specialty, between 2012 and 2016. The choice set is given by the set of hospital referral regions in the United States as defined by Dartmouth Atlas.

A.2.5 Evolution of Welfare Differences between Primary Care Physicians and Specialists

The wage differential as well as the reimbursement differential between primary care physicians and specialists has been increasing over time. The differences are calculated using log hourly wages for specialists and primary care and the average is taken across all hospital referral regions for every year. The increase in the wage gap between primary care physicians and specialists between 2012 and 2016 has been equal to 0.037 log units (over $1/hour).

I calculate the welfare as follows:

Table A.4: Welfare Decomposition: All Locations

	ΔCompensation	ΔCompensation, Rent, Amenities	ΔAll
2012	1.36	1.36	1.36
2013	1.36	1.26	2.11
2014	1.24	1.09	2.31
2015	1.31	1.24	2.09
2016	1.40	1.51	1.93
Δ2016-2012	0.04	0.15	0.57
$\frac{\Delta\text{Welfare Gap}}{\Delta\text{Wage Gap}}$	1.10	4.16	15.48

Notes: These results come from the welfare analysis described in the paper. I measure the physician's willingness to pay (in log wages) to pick his first choice counterfactual location instead of his first choice location in 2012. I then analyze the expected utility change driven by different factors and calculate the average welfare gap between specialists and primary care physicians in the different environments. The choice set is given by the set of hospital referral regions in the United States as defined by Dartmouth Atlas. I account for the initial wage gap, reimbursement gap, and reimbursement rate gap observed in the data. The initial wage gap between primary care physicians and specialists in the past 5 years is equal to 0.037 log units (over $1/hour). More details on the estimation are available in the text.

$$\text{Welfare}_t^k = \log\left(\sum_j \exp\left(\delta_{jt}^{k,\,\ell} + \beta_j^{k,\,\ell} x_{ijt} + \epsilon_{ijt}\right)\right) \tag{42}$$

I then measure the physician's willingness to pay (in log wages) to pick her first choice counterfactual hospital referral region instead of her first choice hospital referral region in 2012. I look at the expected utility change driven by different factors and calculate the average welfare gap between specialists and primary care physicians resulting from such differences.

In particular the change in the welfare gap between specialists and primary care physicians in year t is given by the following:

$$\Delta\text{Welfare Gap}_t = \left(\text{Welfare}_t^{SP} - \text{Welfare}_t^{PC}\right) - \left(\text{Welfare}_{2012}^{SP} - \text{Welfare}_{2012}^{PC}\right) + \left(w_{2012}^{SP} - w_{2012}^{PC}\right) + \left(reimb_{2012}^{SP} - reimb_{2012}^{PC}\right) \tag{43}$$

The results from the four different welfare gap analyses (total compensation, total compensation with rents and amenities, and all factors) are reported in Table A.4. I find that only considering the wage gap represents well the welfare gap caused by differences in the total compensation between specialists and primary care physicians. However, once everything else is taken into account (rent, amenities), the wage gap alone only captures a fifteenth of the welfare gap between the two physician groups.

I then calculate the same gap by analyzing the differences along the urbanity index, differentiating among big cities, small cities, and rural areas. The current wage gap is highest in big cities, at 0.097, as expected, since physicians are compensated more in places with lower amenities.

Once the geographical differences are taken into consideration, I find that considering only the compensation dramatically understates the welfare gap between the two types of physicians rurally. Compensation alone is not a good proxy of the welfare gap due to the higher wages that make up for lost amenities. All results from the decomposition by geographical area are reported in Tables A.5 through A.7.

Table A.5: Welfare Decomposition by Location Urbanity: Compensation

ΔCompensation	All	City	Small City	Rural
2012	1.36	1.16	1.41	1.43
2013	1.36	1.18	1.38	1.38
2014	1.24	1.11	1.36	1.10
2015	1.31	1.17	1.36	1.29
2016	1.40	1.19	1.47	1.40
Δ2016-2012	0.04	0.03	0.06	-0.03
$\frac{\Delta \text{Welfare Gap}}{\Delta \text{Wage Gap}}$	1.10	0.30	3.00	-1.11

Notes: These results come from the welfare analysis described in the paper. I measure the physician's willingness to pay (in log wages) to pick his first choice counterfactual location instead of his first choice location in 2012. I then analyze the expected utility change driven by changes in compensation and calculate the average welfare gap between specialists and primary care physicians along the urbanity index. The choice set is given by the set of hospital referral regions in the United States as defined by Dartmouth Atlas. I account for the initial wage gap, reimbursement gap, and reimbursement rate gap observed in the data. The initial wage gap between primary care physicians and specialists in the past 5 years is equal to 0.097 log units in cities, 0.020 log units in small cities, and 0.028 log units in rural areas. More details on the estimation are available in the text.

Table A.6: Welfare Decomposition by Location Urbanity: Compensation, Rent, Amenities

ΔCompensation, Rent, Amenities	All	City	Small City	Rural
2012	1.36	1.16	1.41	1.43
2013	1.26	1.12	1.28	1.24
2014	1.09	1.07	1.43	0.84
2015	1.24	1.28	1.28	1.30
2016	1.51	1.39	1.53	1.38
Δ2016-2012	0.15	0.22	0.12	-0.05
$\frac{\Delta \text{Welfare Gap}}{\Delta \text{Wage Gap}}$	4.16	2.29	6.05	-1.78

Notes: These results come from the welfare analysis described in the paper. I measure the physician's willingness to pay (in log wages) to pick his first choice counterfactual location instead of his first choice location in 2012. I then analyze the expected utility change driven by changes in compensation, rent, and amenities and calculate the average welfare gap between specialists and primary care physicians along the urbanity index. The choice set is given by the set of hospital referral regions in the United States as defined by Dartmouth Atlas. I account for the initial wage gap, reimbursement gap, and reimbursement rate gap observed in the data. The initial wage gap between primary care physicians and specialists in the past 5 years is equal to 0.097 log units in cities, 0.020 log units in small cities, and 0.028 log units in rural areas. More details on the estimation are available in the text.

A.3 Unobserved Amenities

From the results obtained paper, I can quickly solve for the implied unobserved preferences for the different job locations. I can then look at whether the unobserved amenities for primary care physicians are related to the unobserved amenities for specialists. Table **??** presents these estimates.

The specialists' utility value of changes in unobserved amenities is positively correlated with the primary care physicians' utility value of changes in unobserved amenities for the same location. The variation explained is about 16 percent. While this is not too high, there are many things that specialists could value more than primary care physicians, such as the presence of hospitals at close proximity, where specialists often have an office and perform procedures.

To see whether the estimations make sense, I rank the locations based on the unobserved amenities for primary care physicians and specialists.

I find that the top locations for unobservables for specialists are: Houston, New York, Miami, Philadelphia,

Table A.7: Welfare Decomposition by Location Urbanity: All Factors

ΔCompensation, Rent, Amenities	All	City	Small City	Rural
2012	1.36	1.16	1.41	1.43
2013	2.11	1.91	2.16	2.19
2014	2.31	2.12	2.37	2.37
2015	2.09	1.89	2.15	2.18
2016	1.93	1.73	1.98	2.03
Δ2016-2012	0.57	0.57	0.57	0.60
$\frac{\Delta \text{Welfare Gap}}{\Delta \text{Wage Gap}}$	15.48	5.87	28.58	21.50

Notes: These results come from the welfare analysis described in the paper. I measure the physician's willingness to pay (in log wages) to pick his first choice counterfactual location instead of his first choice location in 2012. I then analyze the expected utility change driven by the intertemporal changes in the different factors and calculate the average welfare gap between specialists and primary care physicians along the urbanity index. The choice set is given by the set of hospital referral regions in the United States as defined by Dartmouth Atlas. I account for the initial wage gap, reimbursement gap, and reimbursement rate gap observed in the data. The initial wage gap between primary care physicians and specialists in the past 5 years is equal to 0.097 log units in cities, 0.020 log units in small cities, and 0.028 log units in rural areas. More details on the estimation are available in the text.

Table A.8: Unobserved Amenities Implied by the Model

	ΔUnobserved Amenities, Primary Care
ΔUnobserved Amenities, Specialists	0.33
Constant	0.26
R^2	0.16

Notes: The regression uses the unobserved characteristics backed up from the model of physician supply. Changes in the specialists' utility value of unobserved characteristics are correlated with changes in the primary care physicians' value for the same location.

Atlanta, Dallas. All of these are big cities with a high concentration of hospitals and are well-known as desirable locations, reinforcing the validity of the results.