NBER WORKING PAPER SERIES

WHO VALUES HUMAN CAPITALISTS' HUMAN CAPITAL? THE EARNINGS AND LABOR SUPPLY OF U.S. PHYSICIANS

Joshua D. Gottlieb
Maria Polyakova
Kevin Rinz
Hugh Shiplett
Victoria Udalova

Working Paper 31469
http://www.nber.org/papers/w31469

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2023

This manuscript is intended to inform interested parties of ongoing research and to encourage discussion. Any views expressed are those of the authors and not those of the U.S. Census Bureau. Access to the data has been approved for this project by the U.S. Census Bureau, CMS, IRS, and SSA because it benefits Census Bureau programs in accordance with Title 13 of the U.S. Code. The Census Bureau's Disclosure Review Board and Disclosure Avoidance Officers have reviewed this data product for unauthorized disclosure of confidential information and have approved the disclosure avoidance practices applied to this release, authorization numbers CBDRB-FY23-0319 and CBDRB-FY2023-CES005-024. We thank Natalia Khoudian, Jacob Morris, Oscar Chan, Ilana Salant, Daniel Sonnenstuhl, and Marissa Hindelang for outstanding research assistance. We thank Doug Staiger, Owen Zidar, Jeffrey Clemens, Molly Schnell, Janet Currie (our discussants), Jay Bhattacharya, Scott Blatte, Stephane Bonhomme, Zarek Brot-Goldberg, Garret Christensen, Wojciech Kopczuk, Neale Mahoney, David Molitor, Sean Nicholson, Adam Oppenheimer, Raffaelle Saggio, Ilya Strebulaev, John Voorheis, and seminar audiences for helpful comments. We also thank CPAs and attorneys for their advice on understanding physician tax returns. Part of this research was conducted while Udalova was visitor at the Stanford Institute for Economic Policy Research, whose support she gratefully acknowledges. We also thank the Enhancing Health Data (EHealth) program at the U.S. Census Bureau, the Becker-Friedman Institute, and SSHRC for providing support. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Who Values Human Capitalists' Human Capital? The Earnings and Labor Supply of U.S. Physicians
Joshua D. Gottlieb, Maria Polyakova, Kevin Rinz, Hugh Shiplett, and Victoria Udalova
NBER Working Paper No. 31469
July 2023
JEL No. I13,I18,J24,J31

ABSTRACT

Is government guiding the invisible hand at the top of the labor market? We use new administrative data to measure physicians' earnings and estimate the influence of healthcare policies on these earnings, physicians' labor supply, and allocation of talent. Combining the administrative registry of U.S.~physicians with tax data, Medicare billing records, and survey responses, we find that physicians' annual earnings average $350,000 and comprise 8.6% of national healthcare spending. The age-earnings profile is steep; business income comprises one-quarter of earnings and is systematically underreported in survey data. There are major differences in earnings across specialties, regions, and firm sizes, with an unusual geographic pattern compared with other workers. We show that health policy has a major impact on the margin: 25% of physician fee revenue driven by Medicare reimbursements accrues to physicians personally. Physicians earn 6% of public money spent on insurance expansions. We find that these policies in turn affect the type and quantity of medical care physicians supply in the short run; retirement timing in the medium run; and earnings affect specialty choice in the long run.

Joshua D. Gottlieb
University of Chicago
Harris School of Public Policy
1307 E. 60th St.
Chicago, IL 60637
and NBER
jgottlieb@uchicago.edu

Maria Polyakova
Center for Health Policy
Stanford School of Medicine
Encina Commons, Room 182
615 Crothers Way
Stanford, CA 94305
and NBER
maria.polyakova@stanford.edu

Kevin Rinz
U.S. Bureau of the Census
4600 Silver Hill Road
Washington, DC 20233
kjrinz@gmail.com

Hugh Shiplett
162 Westmorland St
Fredericton, NB E3B3L7
Canada
hugh.shiplett@unb.ca

Victoria Udalova
U.S. Census Bureau
4600 Silver Hill Road
Washingon, DC 20233
victoria.m.udalova@census.gov

An online appendix is available at http://www.nber.org/data-appendix/w31469

The [healthcare] industry is not very good at promoting health, but it excels at promoting wealth among healthcare providers, including some successful private physicians who operate extremely profitable practices.

(Case and Deaton, 2020)

My hand surgeon should have been paid $4.5 billion for fixing my broken wrist, not $1,000.

(Crawford, 2019)

The medical profession has changed substantially since Friedman and Kuznets (1945) emphasized the importance of entry barriers. Most significantly, far more patients are now covered by health insurance under contracts that restrict physicians' payment rates. What determines a profession's earnings when its output faces regulated prices yet potential entrants face high barriers? We use a new U.S. tax data linkage to analyze earnings and labor supply of the highly skilled and highly regulated medical profession. We show how payment policies shape physicians' earnings and how these incomes in turn affect labor supply.

This occupation merits detailed study as a large share of high earners in the United States are physicians,[1] so documenting their labor market rewards is central to understanding how society values and allocates top talent. The allocation of human capital across different activities is key to how a sector or an entire economy functions (Murphy, Shleifer and Vishny, 1991), and the government's pronounced role in the physician labor market may give it unique power to drive talent allocation of these quintessential "human capitalists" (Smith et al., 2019). We examine this role theoretically, with reduced-form analyses, and with an empirical model of specialty choice.

We begin with a conceptual model of physicians' specialty choice and labor supply in a world with restricted entry, regulated prices, and heterogeneous ability. The model allows

[1] Physicians are the most common occupation in the top percentile of the income distribution (Gottlieb et al., 2023).

us to interpret the impacts of government policy and understand how the high incomes of physicians in exclusive specialties interact with the costs these physicians endured during training. The model shows how limited entry for some specialties, combined with regulated pricing, generates powerful interests to protect these *ex post* rents. When entry is restricted, higher government payments are extra valuable for incumbents. The model illustrates the subtle ways government reimbursement policy shapes the ability distribution across specialties when entry is restricted: sorting is key, rather than the number of entrants, and is nonmonotonic in ability. This framework guides our estimation choices and interpretation.

Our empirical work begins with novel descriptive facts essential to understanding physicians' labor markets. Physician earnings comprise 8.6 percent of total healthcare spending. While this market is of longstanding academic interest (Friedman and Kuznets, 1945; Feldstein, 1970; Fuchs and Kramer, 1973; Sloan, 1975), the modern literature has been hamstrung by measurement challenges that obscure even basic facts.[2] We document the level and composition of physician earnings, how earnings evolve over time, and the pronounced differences across geography and specialty.

We use two-way fixed effects to disentangle the contributions of individual and geographic factors. We find an unusual geographic pattern: rural areas have positive location effects and there is negative physician-location sorting. That is, smaller markets attract lower-earning physicians but boost their earnings. This differs from lawyers, whose pattern we examine separately, and other workers and industries. One natural hypothesis for this unusual pattern is the government subsidies that permeate this market. The tremendous demand increase spurred by insurance (Finkelstein, 2007) and centrally set reimbursements for healthcare services may increase physicians' earnings in rural markets relative to other occupations.

We use two types of policy variation—changes to insurance coverage and to payment rates

[2]This literature (*e.g.* Baker, 1996; Nicholson and Souleles, 2001; Bhattacharya, 2005; Vaughn et al., 2010; Nicholson and Propper, 2011; Esteves-Sorenson and Snyder, 2012; Chen and Chevalier, 2012; Jagsi et al., 2012; Seabury et al., 2013; Altonji and Zhong, 2021; LoSasso et al., 2020; Gottlieb et al., 2023) has relied on survey data and faced measurement challenges, such as top-coding and complicated income structures. Our data overcome many (though not all) of these issues and allow us to newly establish basic facts about U.S. physicians' earnings. Appendix A presents a survey on the public's beliefs about physician earnings.

per service—to distinguish government's influence from other differences across markets that affect physicians' earnings and labor supply. The government's influence is dramatic. When Medicare reimbursements change, one quarter of marginal revenue from public and private insurance flows into physician earnings. When the ACA permanently increased insurance coverage, 6% of public spending accrued to physicians personally.

It is often hard to know whether, beyond accruing rents (Kline et al., 2019; Bertrand, 2009), top incomes play an important allocative role. To answer this, we study labor supply responses to the same insurance coverage and payment policy changes using income tax, Medicare billing, and specialty choice data. We find positive labor supply responses, with a procedure-level short-run supply elasticity of 0.4. In the medium-run, doctors who are past the lifetime earnings peak delay retirement when they experience positive earnings shocks.

Perhaps the most important dimension of labor supply in the long run is talent allocation across specialties. Specialty choice is particularly important because it is sticky; once physicians decide early in their careers, they are unlikely to change specialty later on, so this initial choice has long-run ramifications. It is also complex due to binding entry restrictions for some specialties: physicians cannot simply enter more lucrative specialties at will.

We thus consider a more subtle prediction: higher-ability physicians, who have more choice of specialty, will displace lower-ability physicians in specialties that become more lucrative. We use data on specialty choice separately by physician ability to estimate the earnings elasticity of specialty choice. Our estimates imply that increasing primary care earnings by 5%, while holding constant the number of available slots and earnings of other specialties, would increase the probability of graduates from top-5 medical schools entering primary care by 4.8%. Increased earnings attract physicians with higher test scores to a specialty while displacing those with lower test scores and less choice.

Taking these results together with policy's sizable impact on earnings, we conclude that government payment rules play a key role in valuing and allocating one of society's most expensive assets: physicians' human capital. Policies subsidizing surgery will increase surgeons' incomes and attract more top talent to surgical specialties, improving surgery for a generation. Subsidizing primary care instead will increase these physicians' incomes and attract top talent to primary care, improving primary care for a generation.

These results are key to understanding equilibrium in the market for physicians. They teach us how policies drive consequential short- and long-run outcomes, and provide a clear agenda for future research. Our results show that policy evaluation in this environment must account for the health impacts, and thus social returns, to physician ability in different specialties—currently unknown parameters. We encourage future work to estimate these in order to determine the welfare impact of talent allocation and hence insurance policies.

1 Model of Earnings, Specialization, and Labor Supply

We start with a model that analyzes physicians' labor supply decisions together with policy instruments that can change earnings. Our setting differs from recent work studying the interaction between policy environments and the allocation of talent in the economy (Bell et al., 2019; Hsieh et al., 2019; Jaimovich and Rebelo, 2017) because of entry restrictions (as in Murphy et al., 1991), and the government's direct influence on payments for different categories of workers in this sector.

1.1 Setup

Consider a unit mass of physicians, indexed by i, each with medicine-specific ability $a_i \geq 0$ and idiosyncratic preferences $\epsilon_i \geq 0$ for specializing. Ability follows a Pareto distribution $F(a) = 1 - (1/a)^\theta$ with shape parameter $\theta \geq 1$. Preference for specializing follows an independent exponential distribution $H(\epsilon) = 1 - \exp(-\epsilon)$.

In period 1, each physician chooses whether to specialize ($s_i = 1$) or be a generalist ($s_i = 0$). In period 2, government arbitrarily chooses reimbursement rates for each specialty,

r_s. In period 3, physicians choose intensive margin labor supply L_{is} and earn y_{is}.[3] A physician's payoff includes utility $V(y_{is}, L_{is})$ over consumption (which equals income) and leisure plus direct (dis)utility from specializing $W(s_i; \epsilon_i, a_i)$:

$$U(L_{is}, s_i; a_i) = V(y_{is}, L_{is}) + W(s_i; \epsilon_i, a_i)$$

where:

$$V(y_{is}, L_{is}) = \ln\left(\frac{y_{is}^{\eta}}{\eta} - \frac{L_{is}^{\psi}}{\psi}\right)$$

$$W(s_i; \epsilon_i, a_i) = \left[\epsilon_i - \gamma \ln\left(\frac{A}{a_i}\right)\right] \cdot \mathbb{1}_{s_i=1}$$

with $\eta \in (0,1)$ and $\psi > 1$. The labor/leisure tradeoff is separable in consumption and leisure, with concave utility over consumption (since $\eta < 1$) and convex disutility of work effort L_{is} (since $\psi > 1$). Income y_{is} that depends on the reimbursement rate and hours worked according to a function that may be convex: $y_{is}(L_{is}) = \frac{r_s L_{is}^{\alpha}}{\alpha}$, with $\alpha \in [1, \psi)$.[4] The direct utility from specializing is normalized to zero if the physician chooses to be a generalist. If she chooses to specialize, the utility has two terms: an idiosyncratic match component ϵ_i and a cost of studying for the specialty, $\gamma \ln\left(\frac{A}{a_i}\right)$. We assume $A > 0$ and $0 < \gamma < \theta$ so high-ability physicians have a comparative advantage in specializing.

1.2 Physician Choices

Physicians make two career choices in this framework: specialty in period 1 and the amount of labor in period 3. We start with the labor-leisure tradeoff. Our assumptions generate

[3] We assume that ability only matters when studying to enter a specialty, but our results would strengthen if ability also impacts earnings directly.

[4] The assumption that only reimbursement rates and labor effort determine earnings is obviously a simplification. One natural generalization would be to introduce additional dimensions of physician heterogeneity. By focusing exclusively on the reimbursement rate r_s, we also ignore private insurance negotiations, employment relationships, and costs of production, which could vary across markets. We will use physician movers to empirically disentangle worker heterogeneity from market-level factors in Section 4. This allows us to determine the income variation across space relevant to physicians' choices—factors akin to r_s—separately from unobserved individual ability.

upward-sloping labor supply in response to the realized reimbursement rate r_s: $L_{is}^* = r_s^{\frac{\eta}{\psi-\alpha\eta}}$. This implies a stronger upward-sloping *income*-reimbursement relationship conditional on specialty choice:

$$y_{is}^* = \frac{1}{\alpha} r_s^{\frac{\psi}{\psi-\alpha\eta}}. \tag{1}$$

The resulting subutility over labor and leisure, taking specialty choice and reimbursements as given, is then $V_s^* = \chi + \varphi \ln(r_s)$ for specialty s.[5] This is more elastic to reimbursement rates than is labor supply—though less elastic than is income, as changing work effort dissipates some of the earnings difference.

In this framework, demand shocks—i.e. a change in reimbursement rates—are not moderated by entry and competition for rents on the supply side. This happens because we do not allow for entry in period 3—a realistic benchmark given the long training and other entry barriers into medicine. So the impact of reimbursements on rents is absorbed fully by the incumbents, who then have an incentive to lobby and desire to restrict disadvantageous changes. Any competition occurs at earlier stages, when doctors choose specialties.

To study specialty choice, we first work back to period 2 in which the government chooses reimbursement rates. From the physician's perspective, government's price setting in period 2 is random. Assume it is drawn from a lognormal distribution: $\ln r_s \sim \mathcal{N}\left(\ln \mu_s, \sigma_s^2\right)$. The expected subutility from consumption and leisure, $\mathbb{E}V_s^*$, can then be written as $\mathbb{E}V_s^* = \chi + \varphi \ln \mu_s$ for specialty s. Doctor i will choose to specialize when the total expected utility from specializing is greater than the expected utility from being a generalist:

$$\chi + \varphi \ln \mu_1 + \epsilon_i - \gamma \ln\left(\frac{A}{a_i}\right) > \chi + \varphi \ln \mu_0 \tag{2}$$

[5]The constants are $\chi = \ln\left(\frac{1}{\alpha^\eta \eta} - \frac{1}{\psi} > 0\right)$ and $\varphi = \frac{\psi\eta}{\psi-\alpha\eta} > 0$.

Given the distribution of ϵ_i, the share of doctors of ability a_i who *prefer* to specialize is:

$$\Pr(u_{i1} > u_{i0}|a_i) = \left(\frac{\mu_1}{\mu_0}\right)^{\varphi} \left(\frac{a_i}{A}\right)^{\gamma}. \tag{3}$$

This assumes that any physician can choose to specialize or become a generalist. We next account for the empirical reality that not all physicians have this choice.

1.3 Equilibrium Specialty Allocation

In practice physicians cannot choose specialties at will, as the number of specialized training slots n offered in any year is typically lower than the latent demand for these slots. We assume the application process for specialized training ranks physicians according to ability, only admitting those above some cutoff, denoted a_m. In practice, this happens through the residency or fellowship match process and the cutoff a_m is endogenous to the set of applicants. Given this rule, the choice expression (3) is only valid for physicians whose ability exceeds the cutoff a_m. To solve for a_m, we integrate the choices given by (3) over the ability distribution of doctors and equate it to n, yielding:

$$a_m \propto \left[\frac{1}{nA^{\gamma}} \left(\frac{\mu_1}{\mu_0}\right)^{\varphi}\right]^{\frac{1}{\theta-\gamma}}. \tag{4}$$

This equation describes the way physician specialty choices achieve equilibrium when slots are fixed.[6] Higher relative wage expectations $\frac{\mu_1}{\mu_0}$ for a specialty induce supply. But with demand inelastic at n, rents must be dissipated somehow. Equation (4) shows two mechanisms by which this can occur. The first is by reducing the nonpecuniary benefits of specializing—captured here as making the entry process costly (higher A). The second is rationing based on a strict application process, modeled here as an increase in the ability threshold for specializing, a_m. This expression is intuitive: the higher relative wages are for specializing,

[6] The constant of proportionality in (4) is $\left[\frac{\theta}{(\theta-\gamma)}\right]^{\frac{1}{\theta-\gamma}}$.

the higher the ability cutoff. The more physicians are allowed to specialize (higher n), the lower the cutoff. The higher the training cost shifter A—i.e. the less attractive is the specialized work—the lower the cutoff.

Empirical Implications

Since we cannot perfectly measure ability—let alone whatever metric is used in the residency matching process—we develop the model's empirical predictions that can be used with limited data. We consider the match process for doctors in a high-ability regime—those who have free choice of specialty—and those with fewer choices.

High-ability physicians. Suppose we do not see a doctor's exact ability a_i, but do have an indicator that this doctor has high ability; specifically, $a_i > \overline{a}$ for some threshold $\overline{a}$. This is analogous to graduating from a top-5 medical school, or a test score interval, as we use in Section 6.1.[7] The share choosing to specialize is:

$$\Pr(u_{i1} > u_{i0} | a_i > \overline{a}) = \frac{1}{A^{\gamma}} \frac{\theta}{\theta - \gamma} \left(\frac{\mu_1}{\mu_0} \right)^{\varphi} \overline{a}^{\gamma}. \tag{5}$$

This share is increasing in relative wages $\frac{\mu_1}{\mu_0}$ and in the strictness of the cutoff, $\overline{a}$. Furthermore, two variables have a positive cross-partial. This has a clear empirical implication: moving up the ability distribution, the earnings elasticity of specialty choice increases.

Low-ability physicians. Now suppose we observe a doctor with low ability, $a_i < \underline{a}$.[8] Accounting for the endogeneity of a_m with respect to other parameters yields:

$$\Pr(u_{i1} > u_{i0} | a_i < \underline{a}) = n - \frac{1}{A^{\gamma}} \frac{\theta}{\theta - \gamma} \left(\frac{\mu_1}{\mu_0} \right)^{\varphi} \underline{a}^{\gamma}. \tag{6}$$

[7]The key assumption is that $\overline{a} > a_m$ for any set of parameters that is empirically relevant.

[8]The case where $\underline{a} < a_m$ is trivial, as these doctors always have zero probability of specializing. So we assume $\underline{a} > a_m$. The share of this group choosing to specialize is a weighted average of (3), for those with $a_i \in [a_m, \underline{a}]$, and zero, for with with $a_i < a_m$.

This expression is intuitive: there are n slots available in total, and the second term parallels equation (5)—the choice share for higher-ability doctors ($a_i > \underline{a}$). The difference is what is left over for the doctors with $a_i < \underline{a}$. This expression is decreasing in relative wages $\frac{\mu_1}{\mu_0}$; higher wages make the specialty more attractive for the higher-ability doctors, leaving fewer slots remaining for the lower-ability ones. Empirically, we should see the share of low-ability doctors decreasing as a specialty's wages increase. The magnitude of this effect is stronger the more negatively selected the group being measured (the lower the value of $\underline{a}$).

1.4 Discussion

This model captures key features of the physician market and guides our empirical analysis of physicians' earnings and labor supply. First, the model demonstrates the importance of estimating pass-through from reimbursements to income; the earnings-reimbursement relationship in (1) is key for subsequent labor supply and specialty choice analysis. If reimbursement changes were competed away or absorbed by intermediaries such as hospital employers (not modeled here), it would be harder for them to influence physicians' labor supply and specialty choice. We therefore estimate the impact of reimbursements on physician earnings and labor supply in Section 5.

Earnings potential serves two key roles in the model: First, it induces physician labor supply—and, with a generalization, other investments in patient care. Second, it allocates physician talent across specialties. This can have a powerful impact on supply even if the number of physicians in a specialty can't adjust, by working instead through changes in the ability distribution. Higher expected earnings imply some combination of higher *ex ante* entry costs and stronger selection of those able to enter. We test this empirically in Section 6.

Strong selection and high entry costs may help to explain the political economy of physician labor markets. Physicians who specialize will have incurred substantial costs to do so, whether through the difficulty of training or foregoing other opportunities commensurate with their high ability. If they subsequently experience lower-than-expected earnings—

r_s is set below expectations or additional specialists are allowed to enter and drive down earnings—the incumbents will be displeased. Those at the margin will have (*ex post*) made the wrong specialty choice; the inframarginal specialists will see significant rent dissipation. It would be natural to expect them to expend resources lobbying to maintain higher reimbursements and restrict entry.

2 Institutional Background, Data, and Measurement

2.1 Career in Medicine

This section describes the standard sequence of medical training and career progression, important background for our data and measurement choices. A career in medicine is competitive and follows a relatively rigid script. Practicing physicians choose specializations early and these are hard to change. Physicians' earnings can be complex and frequently include both wages and business income.

Medicine is a professional degree in the U.S. A high school student who wants to become a physician must first complete an undergraduate degree and then earn an MD from one of 158 medical schools. Around 50,000 students apply to U.S. MD-granting medical schools annually and around 45% are admitted (AAMC, 2022). The top-ranked schools are highly competitive; Stanford admits 2.2% of applicants and Harvard reports an average undergraduate GPA of 3.9. Halfway through their (usually) four years of medical school, students take the first standardized test required for the U.S. medical license, USMLE Step 1.

To practice medicine, MD recipients must next complete a residency in a specific specialty. Residency programs take several years, but vary substantially in their competitiveness and length.[9] Primary care is typically less competitive and shorter than the more specialized programs. After completing residency, physicians can begin to work in private practice,

[9] Overall, 5,313 residency programs offered 36,277 positions in 2022. 19,902 U.S. MD graduates applied and 93% received an offer (NRMP, 2022).

small groups, or larger organizations, or complete further "fellowship" training.

The earnings structure of independently practicing physicians can be classified into three broad models. One extreme is physicians who only earn wage or salary income reported on Form W-2. This is common in larger organizations such as academic medical centers. The second model is a sole proprietorship, generating income that only appears on Schedule C of the physician's personal tax return ("Profit or Loss from Business, Sole Proprietorship"). The third model involves a pass-through entity, usually an S-corporation or a partnership. A medical practice organized as an S-corporation pays physicians a market wage, reported on W-2, plus a share of profits that remain after all practice expenses. The exact legal structure affects the tax liability and the profit-sharing incentives within the practice.

2.2 Data Sources

Our primary source is the universe of individual federal income tax returns from 2005 to 2017 merged with an administrative registry of all healthcare providers in the U.S. We augment these with additional administrative and survey datasets.[10]

Tax Data. We use an extract from federal income tax data containing the universe of individual tax returns for tax years 2005 through 2017. We augment individual returns with third-party information returns, notably Forms W-2 and 1099-SSA. Form W-2 reports wage earnings for each filer in the tax unit and includes the Employer Identification Number (EIN) for those physicians who had any W-2 income.[11] We inflation-adjust all monetary values to 2017 dollars using the CPI-U deflator from the BLS and replace missing records with \$0 if the person filed taxes. Tax data also include the state and county of residence.

Physician Registry. We merge tax data with the administrative registry of physicians (the National Plan and Provider Enumeration System, or NPPES) using the Census Bu-

[10]Details of data sources are provided in Appendix B.1.

[11]For ease of exposition, we refer to the tax unit reported in the EIN on Form W-2 as "firm" throughout.

reau's Protected Identification Key (PIK)-based data linkage infrastructure. NPPES lists all physicians and their specialty,[12] which we augment with medical school name, the school's U.S. News and World Report medical school ranking, and the physician's graduation year.

Demographic Data. We obtain date of birth, date of death if applicable, sex, and citizenship status from the Census Bureau's version of the Social Security Administration's Numerical Identification database (Census Numident, as described in, *e.g.*, Bailey et al., 2020; Polyakova et al., 2021). We infer marital status from the tax filing status.

American Community Survey. Using PIK-based linkages, we add responses from the restricted-use version of the American Community Survey (ACS) for those physicians whose household was surveyed by ACS between 2005 and 2017.[13] This provides self-reported earnings and work hours. We also use ACS to construct a sample of lawyers for comparison.

Medicare Data. We add data on treatments physicians provide to Medicare patients. Since 2012, CMS has released data with the list of services performed, the number of times each service was offered, and additional detail by physician. We further add data on Medicare reimbursement rates for each service-year.[14]

2.3 Income and Retirement Definitions

Physician incomes come through diverse and changing mechanisms. This mishmash of sources makes it particularly challenging to study physician earnings and highlights the advantage of using tax rather than survey data to measure income. We construct four measures of income in the tax data: individual total income; individual total wage income

[12]Physician specialty is defined at varying levels of detail; in this paper, we primarily use Medicare's specialty codes or broader aggregates we define in Appendix B.1.

[13]Restricted-use ACS has finer geographic detail, less income top-coding, and a larger sample than the public-use version.

[14]As Section 5.1 describes, Medicare pricing is based on service-specific (and whether the service is provided in a facility or not) Relative Value Units (RVUs), which are recorded in CMS Physician Fee Schedule files.

including any pre-tax deferrals to retirement plans or alike;[15] individual total business income; and Adjusted Gross Income (AGI) at the household level. We define retirement as the year in which an individual who is older than 40 first receives form 1099-SSA, "Social Security Benefit Statement." Details are in Appendix B.2.

2.4 Sample Definitions

Our main sample is a panel of physician-year observations for 2005 to 2017 for physicians aged 20 to 70. This results in 11.6 million physician-year observations for 965,000 unique physicians in our main sample, 848,000 of whom are observed in the 2017 cross-section (Table 1). In many of our analyses we also use two age-based sub-samples: a peak earnings sample of ages 40 to 55 and high retirement risk sample of ages 56 to 70 (350,000 and 287,000 physicians, respectively, in 2017).[16]

3 How Much Do U.S. Physicians Earn?

3.1 Basic Facts

Table 1 reports summary statistics for the full sample (column 1), year 2017 cross-section (column 2), and two age-based sub-samples of this cross-section (columns 3 and 4). The average physician in 2017 earns $243,400 in wages and $350,000 in total individual income. Income is right-skewed; median total individual income is $265,000. One third of physicians have business income exceeding $25,000. At the tax unit level, 24% of physicians are in the top percentile of the national income distribution and median AGI is $325,500. Real earnings of physicians grew by 1% annually over the time period we consider (Appendix B.3).

Physicians in aggregate earn $297 billion in pre-tax dollars measured by total individual income, or 8.6% of total U.S. healthcare spending in 2017 (CMS, 2019). Put differently,

[15] We include deferred contributions into wages and subtract likely deferred account withdrawals from total individual income. The idea is to record earnings in the year they are earned, not consumed.

[16] All numbers in the manuscript are rounded according to U.S. Census Bureau disclosure protocols.

out of \$10,611 that an average American spent on healthcare in 2017, physicians earned \$913. While billing for physicians' clinical services comprises one-fifth of spending, less than half of this amount is physicians' own pay.[17] Subtracting individual income tax payments at a rate of 30% implies that physicians' total after-tax earnings is closer to 6% of total U.S. healthcare spending, or 1% of GDP. This puts an upper bound on the magnitudes at stake in policy discussions that suggest reducing physician incomes in order to lower U.S. healthcare spending (*e.g.* Baker, 2017).

Table 1 reports additional characteristics of physicians and their work environments, including specialty, firm size, work hours, and medical school characteristics. We discuss these factors in Appendix B.3.

Earnings Variation. Figure E.1 illustrates the substantial heterogeneity behind average earnings. More than 25% of physicians in 2017 earn above \$425,000, and the top 1% of physicians earns above \$1.7 million. Table E.2 asks what share of this variation relates to observable characteristics. We run a series of regressions of physician earnings on changing covariates. We first add basic demographics—age, sex, marital status, and whether the individual was ever a non-U.S. citizen. Age accounts for 14% of the variation. Conditional on age, adding other demographics brings R^2 up to 0.19. Women earn 30% less than men. We then consider the explanatory power of covariates that physicians have (at least some) control over throughout their careers—attending a top-5 medical school, specialty, location (commuting zone), size of practice, and presence of business earnings. Specialty and firm size explain (statistically) substantial shares of earnings variation. Physicians who graduate from the very top schools have 12% higher income than others, but this effect operates almost fully through access to specialties. Together, pre-determined demographics and career "choice" variables explain up to 37% of the observed variation.

These results highlight two facts that guide our subsequent analyses. First, age, the

[17] This distinction is a major limitation of previous studies that use health record or claims data to infer something about physicians' own earnings, such as the gender pay gap (*e.g.* Ganguli et al., 2020).

presence of business income, firm size, and specialty appear to play key roles (statistically) in explaining earnings. We flesh out the specific patterns along these dimensions next. Second, conditional on all observables, almost two-thirds of the variation remains unexplained. We come back to this fact in Section 4, where we use non-parametric two-way fixed effects (TWFE) to quantify the roles of observed and unobserved individual-specific and market-specific factors. The TWFE analysis can only address some dimensions, since fixed characteristics (most importantly, specialty) are part of the individual-specific component. Section 6 thus formalizes specialty choice as a central labor supply decision for physicians, taking our theoretical model to the data.

Age Profile. Figure 1A plots individual total income by five-year age groups in 2017. The earnings profile is very steep. Physicians earn around $60,000 on average in their late twenties, while they are still in training. This escalates rapidly to an average of more than $185,000 in the early thirties, and peaks at around $425,000 at age 50. Work hours begin to fall and the probability of retirement starts rising at age 60 (Figures E.4A and E.4B), but earnings remain close to $270,000 into the late 60s. This age pattern motivates our focus on income during ages 40 to 55 in much of subsequent analyses, as that age interval reflects a physician's maximum potential earnings.

Figure 1B shows that the growth in earnings during the highest-earning ages occurs through business income. Average wages are almost flat at around $285,000 at ages 40 to 55, while business income (along with the probability of filing Schedule C; see Figure E.4C) grows steadily and accounts for nearly a quarter of earnings at age 50.

Administrative vs. Survey Data. Figure 1C uses the subset of physicians who responded to the 2017 ACS to highlight the differences between survey and administrative data. Total individual income is measured to be substantially higher in tax data than the ACS—for the exact same individuals. The difference is especially large at the career peak. During physicians' most productive years, the ACS estimates are about $140,000 lower, or one-third of the administrative data mean. Tax-based earnings grow much more rapidly during the highest-growth ages. The difference between the two measures is driven by the extensive margin underreporting of business income in the survey data (Appendix B.4)—a crucial part of physicians' earnings, as discussed above.

Firm Size. Figure 1D shows the relationship between earnings and firm size among 40- to 55-year-old physicians. We see a pronounced non-monotonicity. Physicians in single-physician EINs have the lowest average earnings of $382,000. Average earnings are highest in firms sizes that correspond to small group practices of 8 to 10 physicians. Moving to larger firms—presumably large physician organizations or hospitals—average earnings decline.

Top Earners Among Physicians. Table 2 examines the long right tail of the physician income distribution, showing how top earners differ from average physicians. We focus on physicians age 40 to 55 in 2017. First, as with the general population, the income gradient is steep for these quintessential "human capitalists." The top 1% of physicians averages $4.0 million in annual earnings, 10 times average annual earnings in the sample and more than twice the average earnings in the top 5%.[18]

Second, business income is crucial for the top earners. 80% of physicians in the top 1% report business income of at least $25,000, compared to 44% in the top half and 35% overall. The share of earnings coming from non-W-2 sources is also substantially higher among top earners: 85% for physicians in the top percentile, but 6% for an average physician.

Third, top earners are 67% more likely than the average physician to attend top-5 medical schools and 38% as likely to work in primary care—motivating our model in Section 1. Top earners are 6 times more likely to be neurosurgeons, which is one of the most human capital investment-intensive specializations.

[18]For each cutoff in the table, mean incomes among physicians above the cutoff are very nearly double the cutoff point itself. This suggests the physician income distribution is close to Pareto with a shape parameter of 2 throughout the top half of the physician distribution, as Gottlieb et al. (2023) assumed previously when relating physician and non-physician top earnings.

Overall, the evidence on top physician earners is consistent with top earners in the economy broadly (Smith et al., 2019). The very top incomes are observed among highly trained physicians who earn business incomes rather than wages.

3.2 The Importance of Specialty

Earnings by Specialty. Earnings vary substantially across specialties (Table E.5). Primary care physicians (PCPs), the most common specialty category, is also the lowest-earning. Average total individual income among 40- to 55-year-old physicians is $201,200 ($198,300 at median) for PCPs, or 50% of the overall sample mean in 2017. The highest earners are procedural specialists and surgeons, whose average individual earnings are about twice those of PCPs. This variation is visible in the probability of being in the top 1% of households nationally, which in 2017 ranges from 16% among primary care physicians to 57% among surgeons. PCPs' earnings rose somewhat faster than other specialties and general national incomes in the decade we consider; in 2005 surgeons had income 2.6 times that of PCPs, and that ratio went down to 2.4 in 2017. In contrast, average earnings of anesthesiologists and radiologists fell over this time period.

Correlates of Specialty Income. Higher earnings could make one specialty more attractive than another, or could represent a compensating differential as in equation (4). While Section 6 formally evaluates income's role in allocating physicians' talent across specialties, here we present descriptive relationships between earnings and specialty characteristics. These suggest that higher incomes indeed make specialties attractive rather than just compensate for disamenities.

We first examine how earnings differences across specialties relate to two key job amenities: working hours and training length. Figure 2A shows the relationship between total individual income (ages 40 to 55) and weekly working hours (based on ACS responses) at the granular Medicare specialty level using all years of our data (2005–2017). Specialties in which physicians report longer work weeks (such as neurosurgeons and cardiac surgeons, at close to 65 hours) have higher incomes. Ten extra hours per week is associated with $195,000 higher annual income (or around $375 per hour). Two notable outliers well above the regression line are dermatology (44 hours) and ophthalmology (48 hours). Family practice, internal medicine, and pediatrics are all below the regression line, suggesting a disproportionate drop in hourly earnings. Figure 2B shows a very strong relationship between average income in specialty and physicians' average length of training.[19] Each extra year of training is associated with $143,000 in extra annual income.

The challenges of medical training extend beyond the length alone. For instance, new residents matching in 2020 report having conducted two to three times as much research during medical school as their counterparts a decade earlier (Ahmed and Adashi, 2023). Panel C shows that the level of research experience—among those who successfully match in a specialty—is positively related to a specialty's income. Research experience among matched physicians is an equilibrium choice, so it is not clear whether to view it as a measure of ability or the specialty's entry costs.[20]

Panels A, B, and C show clear relationships, but also a fair amount of variation around the regression lines ($R^2 = 0.36$ in Panel A, $R^2 = 0.54$ in Panel B, and $R^2 = 0.64$ in Panel C). Any specialty above the regression lines must either have compensating differentials for unobserved job characteristics (such as flexibility, time on call, liability risk, or type of work) or be more attractive to potential entrants.

To distinguish between these explanations, we examine labor supply given the bundle of earnings, training, and hours that each specialty offers. Given the presence of entry restrictions, the number of physicians in a specialty is not an appropriate measure of labor

[19] Although training is largely standardized within a specialty, there is variation across programs and across individuals. To systematically determine each specialty's *actual* average training length, we develop a method to estimate it empirically using the tax data. Appendix B.2 provides details. Our measure ranges from 3.7 years for family practice to 7.2 years for cardiac surgery.

[20] In the language of equation (4) in the theoretical model, demanding more research experience would be a higher A, while selecting a higher caliber of medical students would be a higher a_m. Either one would bring the model into equilibrium as a specialty's relative income $\frac{\mu_1}{\mu_0}$ increases.

supply. Instead, we look at *who* enters each specialty. Residency and fellowship programs generally prefer domestic MD graduates to other applicants. So each specialty's share of entrants from U.S. MD programs is a coarse metric of the specialty's appeal to incoming physicians. Figure 2D relates this share to the unexplained part of specialty earnings.[21] We residualize both the share of U.S. MD-trained physicians and specialty mean income with respect to training duration and work hours. We then plot the residualized U.S.-trained share against residualized income (with sample means of each variable added to the residuals). We observe a clear upward slope. Conditional on hours and training, a specialty with $100,000 higher peak earnings tends to have a 7 percentage point higher share of U.S. MD graduates. This suggests that income above the regression lines in Panels A and B is largely an attractive feature of a specialty rather than a compensating differential. In Section 6 we move beyond this descriptive relationship and estimate a formal model of specialty choice.

4 Sources of Variation in Physician Earnings

The geographic pattern of physician earnings is striking. Figure 3A shows average earnings for physicians aged 40 to 55 by state. The pattern is unusual: physicians incomes are not highest on the coasts, as they are for lawyers (shown in Panel B) and the broader economy.

This section unpacks this pattern. In Section 4.1, we use event studies to implement the movers strategy of Finkelstein et al. (2016), Molitor (2018), and others, to determine the causal importance of location. We then delve into a finer decomposition of place- and person-specific factors based on Card et al. (2021) in Section 4.2, and describe the characteristics of high-earning locations in Section 4.3. The place effects for physician earnings are indeed unique, with *negative* sorting between people and places.

These findings suggest that the substantial policy interventions in this market could be

[21] Panels C and D use National Residency Match Program (NRMP) specialty definitions. NRMP data allows us to distinguish between the number of U.S. MD graduating seniors and other applicants who match to a specialty. Non-U.S. MD graduates are primarily graduates of international medical schools, but also include graduates of U.S. DO programs.

important for the geographic patterns, and motivate us to specifically examine the role of government payments in Section 5. Section 5.3 then puts these pieces together to understand the variation we document here.

4.1 Place vs. Physician Factors: Event Study

Empirical Approach. We use physician movers to ask if location matters for earnings. For each physician i who moves across commuting zones,[22] we denote the difference between average log incomes among physicians in i's origin CZ c and destination CZ c' by $\Delta\overline{\ln y}_{(c,c')}$. We then estimate the following regression using data from twelve years around the move:

$$\ln y_{it} = \alpha_i + \sum_{\tau \neq 0} \beta_\tau \times \mathbb{1}_\tau \times \Delta\overline{\ln y}_{(c,c')} + \theta_{a(i,t)} + \lambda_\tau + \varepsilon_{it}, \tag{7}$$

where $\ln y_{it}$ denotes physician i's annual log individual income. This is a dynamic, parametric event study specification, which yields coefficients $\hat{\beta}_\tau$ on the income change for each year τ relative to the year of move ($\tau = 0$). Under the standard assumption that shocks ε_{it} are conditionally mean-independent of location causal effects, the post-move coefficients can be interpreted as the share of the geographic income differences due to place rather than person. Although standard, this assumption cannot be taken for granted, so we use the pre-move $\hat{\beta}_\tau$ coefficients to investigate it. We control for physician (α_i), physician age ($\theta_{a(i,t)}$), and relative time (λ_τ) fixed effects.[23]

Results. Figure 4 shows that location drives a large share of earnings. The estimates of $\hat{\beta}_\tau$ show a sharp change in income at the time of the move and no differential trends in income preceding the move. The point estimates suggest that movers' incomes shift by 70% of the difference between mean incomes in origin and destination. This estimate is even higher

[22] Our sample is all 40- to 55-year-old physicians who changed their commuting zone (CZ) residence exactly once between 2005 and 2017. We use CZs to capture both within- and cross-state variation; unadjusted CZ average incomes are shown in Figure E.5.

[23] Calendar year fixed effects are implicitly included as a linear combination of the other fixed effects.

within location-by-specialty.[24] Having established that location influences earnings, we next examine the patterns of these locations' effects and how physicians sort across them.

4.2 Place vs. Physician Factors: Variance Decomposition

Empirical Approach. We use a two-way fixed effects model to decompose place (c) and person (i) contributions to individual earnings. Year t earnings are:

$$\ln y_{it} = \alpha_i + \psi_{c(i,t)} + \theta_{a(i,t)} + \lambda_\tau + \varepsilon_{it}, \tag{8}$$

in which α_i is the individual component, ψ_c is the location (commuting zone) component, ε_{it} is a mean-independent person-time residual, and some specifications include fixed effects for age, $\theta_{a(i,t)}$, and for time relative to the year of move, λ_τ. Moves must be independent of the shocks ε_{it}, and the lack of pre-trends in Section 4.1 provides some justification for this assumption. Limited mobility bias could plague a naive variance decomposition, so we implement the Andrews et al. (2008) homoskedastic correction, and the Kline et al. (2020) heteroskedastic correction along with a direct fixed effects estimation (Abowd et al., 1999).[25]

Results. Figure 5A shows the key results. The first three bars show the estimated variance of location effects, $\mathrm{Var}(\psi_c)$, using standard fixed effects estimation, the homoskedastic correction, and the heteroskedastic correction. The next three bars report estimates of how physicians sort across space, $2\mathrm{Cov}(\alpha_i, \psi_c)$, for the same three methods. All three show pronounced negative sorting. The magnitude of sorting is substantial relative to that of the location effects themselves; the ratio of covariance to variance is between 0.6 and 0.8 (Table E.6). Column (4) shows that the result is stable when adding time-varying controls.

[24] Figure E.6 shows analogous event study graphs for four subsamples of physicians: specialists, PCPs, and physicians who graduated from ranked and from unranked medical schools. For each sample, we construct the income difference using physicians in the same category. Looking within specialty leads to coefficients meaningfully larger than the overall average, with the point estimates closer to 0.85. This suggests that the key driver of variation is the interaction of specialty and location, and specialty earnings have different geographic patterns.

[25] We use the Bonhomme et al. (2023) `PyTwoWay` package to implement and describe all of these estimators.

Panel B presents analogous estimates for lawyers, another highly-educated occupation, but with a differently structured labor market. We again initially find a negative covariance when using the standard fixed effects estimator, but in this case the limited mobility bias corrections reverse the sign.[26] The magnitude of the corrected covariance is in the same ballpark as for physicians, but with the opposite sign. So our data and procedures do yield the expected positive sorting, consistent with Card et al. (2021) and the broader literature on worker-firm matching—when the data support it. Physicians' pattern is unique.[27]

Importance of Sorting for Geographic Patterns

What does this sorting mean for the overall pattern of earnings across space? We follow Card et al. (2021) and address this question by aggregating equation (8) to the CZ level:

$$\overline{\ln y}_c = \overline{\alpha}_c + \psi_c + \beta\overline{X}_c. \tag{9}$$

This decomposes area-level average log earnings $\overline{\ln y}_c$ into a location effect ψ_c, the *average* person effect among physicians in the location $\overline{\alpha}_c$, and (in some specifications) additional controls.[28] The variance decomposition of (9) reveals what share of variation in areas' average incomes come from the places themselves, the composition of workers, and sorting of those workers across locations.

Figure 5D shows the results by relating the estimates of $\overline{\alpha}_c$ and ψ_c in a binned scatterplot.

[26]The similarity of results between the two correction methods is likely because the `PyTwoWay` package collapses observations at worker-location spell level, which addresses the serial correlation that plagues naive fixed effects estimation. Comparing the lawyer and physician samples, the latter is an order of magnitude larger because we identify physicians with administrative data and lawyers from the ACS sample (see Appendix B.1). The trace of the matrix governing the TWFE bias is an order of magnitude larger for lawyers, which is why the corrections have such an impact for them but little for physicians.

[27]Two further analyses provide evidence that this negative covariance is not an artifact of limited mobility bias. First, we conduct a simple split sample estimation and obtain results extremely similar to those reported here. Second, we estimate a parametric alternative to two-way fixed effects: we use a linear regression to adjust raw earnings for various individual-level *observables*, but *not* individual fixed effects, and take each CZ's fixed effect conditional on these observables. We then correlate these CZ effects with the individual effects from (8). As we add additional covariates, the correlation of individual fixed effects with the conditional CZ effects becomes increasingly negative, trending towards the pattern in Figure 5C.

[28]Including or excluding age fixed effects and relative time fixed effects has little impact on the estimates of location variance and sorting (Table E.7).

The sorting remains negative. The last column of Table E.6 reports the magnitude. The relative magnitude of the sorting effect increases to around 1.2 times the variance of location effects, compared with around two-thirds of the location variance when estimated at the individual level. To benchmark the magnitude, the covariance of CZ-by-industry effects with person effects in Card et al. (2021) explains 1.8 times the magnitude of the CZ-industry effects.[29] The relative magnitude of our sorting is 1.2 times that of location effects, but with the opposite sign.

4.3 Correlates of Fixed Effects

We explore the economics of these places by projecting the place fixed effects on observable characteristics.[30] For a series of location characteristics, Figure 6 shows two correlations: that between the characteristic and our estimated place fixed effect, and that between the characteristic and the raw mean log earnings of the location.

Measures of the location's general economic strength tend to be uncorrelated, or have a slight positive relationship, with the location's raw physician earnings. This pattern holds whether measuring the region's economic strength with income, education, real estate prices, or population size. Life expectancy is slightly negatively correlated with both earnings measures, though the movers-based treatment effect on mortality (Finkelstein et al., 2021) is statistically unrelated. Panel B relates both physicians' and lawyers' place fixed effects to local average income. The difference is striking: lawyers' fixed effect is nearly uncorrelated with local income, while the relationship for physicians is negative and precise.[31]

A few economic forces could generate this pattern. First, physicians could be fundamentally more productive in low-income places, though this contradicts empirical evidence on

[29] Card et al. (2021) consider CZ-by-industry while we consider location effects for one occupation. Our data differ from the LEHD Card et al. (2021) use in that we include the self-employed and non-wage income.

[30] We use estimates for commuting zone fixed effects based on (8) with the full set of controls, but without the limited mobility bias corrections that—as we have shown above—do not change the baseline sorting pattern among physicians. The TWFE analysis in equation (8) also yields fixed effects for each individual physician. Their patterns are similar to the raw physician descriptives discussed in Section 3, so we do not present them further.

[31] Table E.8 reports regression coefficients and standard errors for all regressions shown in Figure 6.

agglomeration in healthcare (Dingel et al., 2023). Second, physicians may face less competition in smaller and lower-income markets and thus be able to charge higher markups to self-paying and privately insured patients. But the magnitude of this force is probably insufficient to explain all of the earnings differences we observe.[32] Third, the income gradient may reflect compensating differentials for skilled workers' preferences to live in higher-income locales. But it is not clear why this would be true for only physicians and not lawyers. Finally, government's major role in the healthcare market, and the complex political economy of this role, could cause outcomes to differ from other industries. Federal and state governments purchase medical services on behalf of lower-income and rural residents, increasing these consumers' effective purchasing power for healthcare relative to other goods or services. Section 5 measures this influence and asks how it affects the geographic earnings gradient.

4.4 Firm Fixed Effects

While intimately linked with the physician's location, the firm at which a physician works may have its own influence on earnings. To explore this, we estimate a firm-worker two-way fixed effects decomposition using a model analogous to (8). The results, shown in Table E.6 Panel C and Figure E.7, are broadly consistent with the location-physician decomposition. Both the raw and bias-corrected covariances between the physician's individual and firm effects are negative, suggesting that physicians' sorting across firms reflects these firms' geographic locations. The negative covariance reinforces the uniqueness of physicians' labor markets (cf. Bonhomme et al., 2023).

[32] Clemens and Gottlieb (2017, Fig. 2) report 20 percent higher private payments in the most concentrated markets compared with the least. Dunn and Shapiro (2014) find that a 10 percent increase in physician market concentration increases prices by 1 percent. In hospital pricing, Cooper et al. (2019) report 12 percent higher prices in monopoly markets than those with at least four competitors. The pattern we observe is also not driven by CZs with very few physicians. Indeed, the negative correlation between CZ FE and median household income is stronger among CZs with more than 10 physicians. Further, CZ FE are only weakly negatively correlated with the number of physicians in each CZ in our data, while we would have expected a pronounced negative correlation driven by small markets if the main underlying mechanism were market power (Bresnahan and Reiss, 1991). In short, these estimates imply that lack of competition is insufficient to drive the scale of differences we estimate in markups across commuting zones.

5 Government Influence on Physician Earnings

We use two empirical strategies to investigate the government's influence on physicians' earnings and short- to medium-run labor supply. We first use short-run price changes, which occur as Medicare adjusts its reimbursements for each procedure. We then examine a persistent demand shock from the expansion of health insurance coverage under the Affordable Care Act (ACA). As this is a permanent demand shock rather than transient variation in prices, it is more informative of longer-run behavioral responses.

5.1 Short-Run Changes: Medicare Rate Adjustments

We first use price shocks to estimate short-run elasticities of income and labor supply. The $900 billion-per-year Medicare program reimburses physicians' professional services based on a fee schedule that defines a Relative Value Unit (RVU) for each medical service. Although this RVU metric is meant to reflect the differences in the time and effort it requires to provide different services, it changes due to periodic reviews and political factors.[33] We use changes in the RVUs assigned to each service to estimate how much Medicare payments influence physician incomes and labor supply.

Do physicians shift their service mix as relative prices change? Given the broad set of price changes Medicare implements each year, we analyze supply responses at the procedure code level. To study physician-level response margins, such as earnings, retirement, and total procedure supply, we use physicians' different baseline service mixes to generate a physician-year reimbursement shock. While each billing code's update is applicable nationally, physicians perform different bundles of services. So each physician is differentially exposed to each year's set of RVU changes. This allows us to use the logic of simulated

[33] We rely on three facts about this system. First, RVUs may vary across time and geography, but not across individual physicians. Second, Medicare's RVU Update Committee (RUC) regularly reviews how many RVUs are assigned to each service. The reviews can be triggered by changes in the service, by Medicare's request, or based on a pre-determined five-year review cycle. Third, the timing of when a particular code, or even codes of which specialty, is reviewed is uncertain. Chan and Dickstein (2019) explain the uncertainty in which specialties will be able to propose code reviews at any given RUC meeting. Appendix C.1 provides more details about the institutional setting and our empirical approach.

instruments to construct each physician's Medicare price exposure.

For each physician i, we first compute the average number of times each service k was performed across all years of our utilization data (2012 to 2017), and denote this by $\overline{q_{i,k}}$. We then multiply this time-invariant quantity measure by the time-varying number of RVUs that Medicare assigns to service k and add them up by physician-year. The result is a series of annual price shocks for each physician, purged of any behavioral response. Mathematically, the composite price for physician i who performs a set K of services in year t is:

$$P_{i,t} = \sum_{k \in K} \overline{q_{i,k}} \times RVU_{k,t}. \tag{10}$$

We label this the *Medicare price instrument*. Figure E.8 shows the distribution of annual shocks to this instrument, $\Delta \ln P_{i,t} = \ln P_{i,t} - \ln P_{i,t-1}$.

We estimate the following empirical relationship to determine how log income, $\ln Y_{i,t}$, responds to the Medicare price instrument:

$$\ln Y_{i,t} = \alpha_i + \beta \ln P_{i,t} + \theta_{a(i,t)} + \eta_{t,s(i)} + \varepsilon_{i,t}. \tag{11}$$

The coefficient of interest, β, is the reduced form elasticity of income Y with respect to the Medicare price instrument.[34] We control for physician fixed effects, α_i, physician age fixed effects, $\theta_{a(i,t)}$, and year-by-specialty fixed effects, $\eta_{t,s(i)}$. The key coefficient β is thus identified off the variation in the composition of services that each individual physicians performs.[35] We run our analysis separately for 40–55 and 56–70-year-old physicians; the former group is prime working age physicians, so we remove trainees with fixed incomes and minimize the mechanical decline in income due to retirements. The latter group is closer to retirement,

[34] To the extent that changes in Medicare fee schedule can trigger changes in private insurers' reimbursement rates, as in Clemens and Gottlieb (2017), our reduced form estimate will capture both the direct and indirect effects of Medicare's reimbursement on physician earnings and labor supply.

[35] The main effects of year and specialty are omitted from $\eta_{t,s(i)}$ as specialty is collinear with the physician fixed effect α_i and year is collinear with the combination of α_i and age $\theta_{a(i,t)}$. To account for the large variability in Medicare billing volumes across providers, we use the average Medicare revenue each physician collected in 2012–2017 as regression weights. We cluster standard errors by Medicare specialty.

allowing us to measure that labor supply margin.

To study Medicare's impact on the supply of medical care, we replace the dependent variable with the log number of RVUs physician i bills in year t, $\ln Q_{i,t}$.[36] Since these RVUs are part of our instrument $P_{i,t}$, we expect a mechanical coefficient of 1 in the absence of any behavioral response. With total RVUs as an outcome, the difference between the coefficient and 1 yields the supply elasticity. A coefficient below 1 indicates income-targeting behavior, while a coefficient above 1 indicates a positive supply elasticity.

To get the elasticity of income to Medicare billing, we estimate an IV setup treating the log Medicare price instrument, $\ln P_{i,t}$, as an instrument for the log total RVUs billed, $\ln Q_{i,t}$, and log income, $\ln Y_{i,t}$, as the dependent variable. To quantify any response via the retirement margin of labor supply, we treat income as the endogenous variable and retirement as the outcome. We estimate both of these IV specifications using two-stage least squares.

To estimate more granular labor supply responses, we count the number of times a physician bills for each code in each year, $q_{i,k,t}$. We directly measure how much a change in the code's own RVU metric, $RVU_{k,t}$, affects this supply:

$$\ln q_{i,k,t} = \alpha_i + \beta RVU_{k,t} + \theta_{a(i,t)} + \eta_{t,s(i)} + \varphi_{\kappa(k)} + \varepsilon_{i,t} \tag{12}$$

where φ_κ is a set of procedure fixed effects.[37] To distinguish between effects on the number of patients treated and the care provided per patient, we also estimate a version in which the dependent variable is the number of unique patients per procedure.

Results

Table 3 and Figure E.9 report the results. A 10% increase in Medicare payment rate leads to a 2.4% increase in professional earnings of 40- to 55-year-old physicians. A substantial

[36] This can be interpreted as the number of services a physician provides, weighted by value. It is formally defined in Appendix C.1, where we also present the estimating equations. For a simpler measure, we also estimate a version that simply counts the number of procedures each physician bills, regardless of value.

[37] The subscript κ is distinct from k because the fixed effects are by HCCPS code, while the unit of observation is at the code-by-place of service level.

component of this change is physicians' behavioral response; a 10% increase in the payment rate leads physicians to bill 4.4% more RVUs. (Recall that we must subtract 1 from the coefficient of 1.437 to get the supply elasticity.) 2SLS estimates imply an elasticity of 0.17 between earnings and prices. This intensive margin labor supply response is a composition of performing 3.9% more unique procedures, and shifting to relatively higher-paid procedures. The procedure-level analysis directly shows that a 10% increase in price per procedure leads physicians to do 3.8% more of this procedure. Nearly the full effect (3.4% out of 3.8%) is driven by performing this procedure on additional patients rather than doing the procedure more frequently for the same number of patients.

Intensive margin responses are broadly similar among 55-to-70-year-old physicians. For this group, we also find a response on the extensive margin. The IV estimate shows that a 10% increase in professional earnings driven by changes in the reimbursement rates leads to a 0.5 percentage point decline in the probability of retirement that year.

To better understand the magnitudes, we convert our earnings estimates into a pass-through—how much do physicians' earnings increase when the government pays one more dollar? Our direct estimates imply that physicians earn \$62 of each \$100 in additional *Medicare* spending. Accounting for Medicare's spillover into private insurance spending (Clemens and Gottlieb, 2017; Clemens et al., 2017), we get a lower pass-through of \$25 for each \$100 of *any* insurance spending.[38] Under either interpretation, pass-through is quite large. Our estimates differ from the modest level of rent-sharing with workers found in response to many other shocks (Card et al., 2018), but are similar to rent-sharing with higher skilled workers who benefit, for example, from patent rents (Kline et al., 2019).

Our results indicate that these marginal earnings have real consequences: paying physicians more increases care provision, as more patients receive better-compensated treatments and physicians delay retirement. We do not observe health outcomes so cannot assess the net social benefits of this marginal spending. The labor supply elasticity of 0.2 to 0.4 that

[38] Appendix C.1 details these calculations.

we estimate is lower than in Clemens and Gottlieb (2014) or Cabral et al. (2021), but is similar to other estimates of compensated wage elasticities (Nicholson and Propper, 2011).

5.2 Persistent Demand Shocks and Long-Run Supply Responses

Short-term price shocks, such as those resulting from RVU changes, are most likely to affect short-run supply margins such as the number of patients treated. To study longer-run responses, we exploit a persistent demand shock: the Affordable Care Act (ACA)'s expansion of health insurance coverage. The ACA increased the insured share of the non-elderly population through two main mechanisms. First, 37 states expanded Medicaid eligibility. Second, means-tested subsidies were offered to individuals for purchasing individual health insurance on newly created health insurance Marketplaces. The ACA became law in 2010, but most of the insurance expansions were implemented in 2014 and 2015. We restrict our analysis to states where the full package of key ACA reforms took place roughly on the same timeline. This leads us to analyze incomes and retirement choices in 24 states that expanded Medicaid in 2014 or early in 2015, coinciding with the rollout of Marketplaces in 2014.[39]

Our identification relies on variation in each county's potential magnitude of insurance coverage expansions. There is more scope for insurance coverage to increase in counties that had a higher share of uninsured population prior to the law's implementation. The share of uninsured under-65 population in 2013, on the eve of ACA expansions, varied from under 10% in some counties of Minnesota to over 30% in some counties of Nevada (Figure E.10). Let $U_{c,2013}$ denote the share of the under-65 population uninsured in county c in year $t = 2013$. We estimate the reduced form impact of insurance expansions on outcomes $Y_{c,t}$ using a county-level panel which covers four years of post-expansion data:

$$Y_{c,t} = \sum_{t=2005, t \neq 2010}^{2017} \beta_t \times \mathbb{1}_t \times U_{c,2013} + \delta_t + \mu_c + \theta g_{c,t} + \epsilon_{c,t}. \tag{13}$$

[39] Appendix C.2 provides more details on our definitions and sources.

where δ_t capture calendar year fixed effects, μ_c are county fixed effects, and $g_{c,t}$ controls for the time-varying shares of physicians of each age in each county. $\hat{\beta}_t$ are the coefficients on year fixed effects interacted with our time-invariant measure of exposure, $U_{c,2013}$, and should be interpreted as relative to 2010, the year in which ACA passed.[40]

To interpret the coefficients $\hat{\beta}_t$ as measuring how much insurance coverage affected outcome $Y_{c,t}$, we need the identifying assumption that growth path of potential outcomes absent ACA rollout over time would have been independent of rates of uninsurance among the non-elderly population in 2013 conditional on covariates. While this parallel trends assumption is not directly testable, the event study specification in (13) allows us to assess whether counties with differential rates of uninsurance in 2013 followed a parallel trend in outcomes prior to 2010. Expectations of future demand may be important for persistent outcomes like retirement or firm structure. These choices may respond to anticipated changes in insurance coverage, and thus to the ACA's passage, rather than realized insurance coverage. In contrast, income is likely to change only once expansions take place and demand increases.

In practice, ACA expansions resulted in only a subset of previously uninsured people taking up insurance. To capture the relationship between $Y_{c,t}$ and the insured population in a county, we estimate the first stage event study that measures how the rate of insurance $I_{c,t}$ in county c in year t changed as a function of the share uninsured in 2013. The specification is the same as in (13), but with $I_{c,t}$ as the outcome. To formally scale the reduced form by the first stage, we collapse the differential time path of treatment effect into the pre- and post-implementation periods.[41] We estimate a 2SLS specification treating the rate of insurance in the under-65 population as the endogenous variable of interest and the rate of uninsured population in 2013 as an instrument that is assumed to only affect outcomes through its effects on insurance.

[40] We use the number of physicians in each county-year as regression weights and cluster standard errors at the county level.

[41] Because of the potential for anticipation effects in long-run decisions, we report two sets of 2SLS specifications: one in Table 4 that defines the pre-period to include all years before insurance expansions began (all years before 2014), and another in Table E.9 that drops the intervening years between the ACA becoming law and its implementation (2011 to 2013). Results are similar.

To measure income and labor supply responses, we estimate the effect of insurance expansions on (log) total individual income and the likelihood of generating extra income through self-employment (as measured by filing Schedule SE) among physicians in their peak earning years (ages 40 to 55). For the population at a higher risk of retirement (ages 56 to 70), we measure the effect of ACA expansion on the probability of retirement.

Results

Figure 7 plots coefficients $\hat{\beta}_t$ for the first stage and the reduced form for individual income (among ages 40 to 55) and retirement (among ages 56 to 70). Table 4 reports the first stage, reduced form, and 2SLS coefficients for all outcomes. The first stage estimate (Column 1 of Table 4) shows that counties with a ten percentage point higher pre-ACA uninsurance rate saw a 4.96 percentage point higher rate of insurance coverage in the post-implementation years, with no noteworthy changes in insurance between 2010 and 2013. As in Finkelstein (2007), expanding insurance coverage reduced pre-existing geographic differences in coverage.

Panel B and columns (2) and (5) of Table 4 show that earnings among physicians aged 40 to 55 grew faster in these more affected areas. We estimate that a ten percentage point higher baseline uninsurance rate leads to 3.9% higher individual earnings four years post expansion, or 2% on average across post-implementation years. Scaling this income effect by the first stage suggests that a ten percentage points higher rate of insurance coverage (a 12% increase over the average sample insurance rate of 85% in 2013) increases physician income by 4.9% across post-implementation years. The elasticity of physicians' earnings to the rate of insurance coverage in the under-65 population is thus 0.41.

Columns (3) and (6) of Table 4 shed some light on how physicians may achieve these changes. The probability that a physician files Schedule SE (self-employment income above $400) increases by 3 percentage points for each 10 percentage point increase in insurance. This proxies for the extensive margin of self-employment and potentially captures increased opportunities to generate side income.

Turning to labor supply in Figure 7C and Table 4 columns (4) and (7), we find that a 10 percentage point higher insurance rate leads to a 1 percentage point decline in retirement probability after the implementation of ACA expansions. This effect emerges after the law is signed rather than after implementation, which we would expect if physicians delay retirement in anticipation of a shift in demand. This evidence suggests that the substitution effect dominates the income effect over the time horizon we consider. Converting the post-implementation estimate to an elasticity, a 12% increase in insurance rate leads to 4.9% higher income and 5.4% less retirement, for a medium-run elasticity of retirement to income of -1.1.[42] This suggests a larger behavioral response in response to a more permanent change in income than we found in response to short-run fluctuations in reimbursement rates.

We use our estimates to ask what share of insurance spending on marginally insured patients goes to physicians—a key issue for the political economy of health insurance.[43] Based on our pooled estimates, 6% of the $110 billion annual spending (CBO, 2016) on the ACA insurance expansion accrued to physicians.[44] Since physicians' baseline share of medical spending is 8.6%, their gain from expansions was slightly less than proportional to their baseline expenditure share.

5.3 Can Government Shape Earnings Variation?

We now consider what our estimates of pass-through imply for the ability of government policies to shape the earnings variation documented in Sections 3 and 4. We first examine whether government policies can explain any of the unusual geographic pattern of physician earnings: higher-earning physicians being in lower-earning areas. We consider reimbursement rates in Medicare—one of the main policy instruments. Medicare adjusts these rates for local

[42] A 2SLS specification that drops the years pre-ACA implementation gives us 3.9% higher income and 6.9% less retirement for a medium-run elasticity of -1.8.

[43] The analogous question among hospitals is well-studied (Garthwaite et al., 2018).

[44] Policy reports suggest that ACA expansion resulted in approximately 5.9 percentage points more people insured among non-elderly in total (Tolbert et al., 2020); the uninsurance rate went down from 16.8% in 2013 to 10.9% in 2015. Applying our elasticity estimate, this expansion led to a 2.4% increase in physician incomes, or $8,400 per physician (2.4% of $350,000). In aggregate for 848,000 physicians in our cross-section, this means $7.1 billion of extra spending, or 6% of the $110 billion in annual spending.

input costs, but the adjustment is incomplete, resulting in effective subsidies to rural areas (GAO, 2022). We conduct an exercise based on Section 5.1 to quantify how this large rural subsidy shapes the geography of physician earnings.

Figure E.11 illustrates this calculation. We consider the elasticity between Medicare's Geographic Adjustment Factor (GAF) for physicians' work—a factor that multiplies Medicare reimbursement rates—and the commuting zone's median household income.[45] This elasticity is a policy choice, and we plot its value on the horizontal axis of Figure E.11. The solid vertical line marks the current empirical estimate of 0.09; the dashed vertical line marks the empirical elasticity of 0.33 between a broader local price index (Diamond and Moretti, 2021) and local incomes. Medicare's differences in reimbursement have a much weaker relationship with geographic differences in income levels than do other prices.

The vertical axis of Figure E.11 shows the elasticity of physician earnings to local income. The empirical values estimated in Section 4 for physicians and lawyers (-0.13 and 0.12, respectively) are indicated with horizontal lines.[46] We use our estimates from Section 5.1 of how changes in Medicare prices affect physicians' earnings to draw a relationship between this policy elasticity and income elasticities. Since the GAF multiplies Medicare RVUs, a change in this index affects physician reimbursements similarly to a change in RVUs.[47] We found that a 10 percent increase in Medicare prices increases earnings by 2.4 percent, including any supply response (see Section 5.1). We draw a red line in Figure E.11 with this slope. This line shows the implied relationship between the elasticity of Medicare reimbursement to local income (x-axis) and the elasticity of physician earnings to local income (y-axis).[48]

The y-axis is a summary metric of the shape of the geographic distribution in physician earnings—an equilibrium outcome—while the x-axis depends on policy choices. If the elas-

[45] We use analysis in MaCurdy et al. (2014) to compute commuting zone level GAF. We use 2016 CZ-level median household income estimates from Chetty et al. (2014).

[46] These estimates are obtained from a regression of the CZ earnings FE from Section 4, separately for physicians and lawyers, on (log) median household income.

[47] The GAF is an approximation of the GPCIs applicable to any particular service; see MaCurdy et al. (2014) for details.

[48] The intercept is set such that the line passes through the empirical elasticities between physician CZ FE and local incomes, and the elasticity between GAF and local incomes.

ticity of Medicare prices to local incomes increased to the level of the elasticity of the general consumption basket (the dashed vertical line farther to the right along the x-axis), our estimates suggest that the geographic variation in physician earnings would look more similar to lawyers'. This finding suggests that Medicare's limited adjustment for local costs alone can explain about a third of the unique pattern of physician earnings, at least as captured by the elasticity of CZ FE to local incomes. That said, this exercise has an important caveat. It uses a short-run estimate of income responses to Medicare payment changes, while the CZ FE-income elasticity is a cross-sectional (and thus long-run) relationship. Longer-run pass-through could be larger (due to private market spillovers; Clemens and Gottlieb, 2017) or smaller (due to entry). In Figure E.11, this would change the slope of the red line. The higher is the pass-through from Medicare payments to earnings, the steeper the red line, and the more important Medicare payments are in driving the geography of earnings.

6 Earnings and Allocation of Talent Across Specialties

While we have seen that government policy drives physician earnings and contemporary labor supply choices, earnings may be even more important if they shape talent allocation over the long run. The model in Section 1 emphasized new physicians' specialty choices. Building on this, we conclude the paper by estimating how income changes the allocation of physician talent across specialties. This analysis takes advantage of the specialty-year panel of earnings we are able to construct using our data linkage.

6.1 Empirical Models of Specialty Choice

Overview. Since new physicians can choose from many specialties, we analyze their decision-making using a discrete choice model. This approach generalizes the binary choice discussed in Section 1, both by allowing for more than two specialties and by incorporating observable and unobservable preference heterogeneity.

We identify the relationship between relative incomes and specialty choice using panel variation in each specialty's earnings, accounting for other time-varying attributes. As Section 1 highlights, only sufficiently high-ability physicians have a free choice of specialty in a world with entry restrictions. For this group, we can interpret the choice-income relationship as reflecting preferences. Other physicians' choices reflect a combination of preferences and rationing. We first develop a discrete choice model of specialty choice, which we estimate on physicians who graduated from top-5 U.S. medical schools.[49] We then estimate a version allowing for heterogeneity by physicians' test scores. For the highest-scoring physicians, this model yields similar results to the version estimated on top graduates. The different income-specialty relationship for lower-scoring physicians quantifies the impact of entry restrictions and allows us to see how changing income shapes the entire distribution of talent.

Discrete-Choice Model Specification. Consider new physician i entering the residency match in year t. Suppose i graduated from a top medical school so has freedom to choose any specialty s out of a set S of nine specialty categories. Physician i chooses specialty s to maximize utility:

$$u_{is} = \alpha M_{st(i)} + \beta_i \phi_{st(i)} + A_s + \varepsilon_{is} \tag{14}$$

where:

$$\varepsilon_{is} \sim \text{Gumbel(0,1)}$$

$$\beta_i \sim N(\beta^0 + \beta^1 D_i, \sigma^2).$$

We denote cohort t's beliefs about hourly earnings in specialty s by $M_{st(i)}$; it is the empirical analogue of the expected log reimbursement rate, μ_s in the theoretical model. Specialty fixed effects A_s account for time-invariant amenity differences across specialties. The vector $\phi_{st(i)}$ captures time-varying features of specialty s we can measure. The random coefficients on

[49] See Appendix B.2 for measurement details.

these observables, β_i, allow for both observed and unobserved heterogeneity in preferences for the specialties' time-varying amenities.[50]

Finally, ε_{is} is the part of individual i's preference for specialty s that the individual knows but the econometrician does not observe. It includes any idiosyncratic beliefs about individual match-specific components of specialty earnings or hours. We assume that this unobserved part of utility is independently and identically distributed with a type I extreme value (standard Gumbel) distribution across all graduates and all specialty choices s.

Identification. A common concern in discrete choice models of consumer demand is that the error term includes a choice-market-specific shock correlated with prices and consumer choices. In our context, this means α will not reflect physicians' true preferences if earnings potential is correlated with the specialty choice constraints they face—which may occur if underlying ability is correlated with both. We address this by limiting the sample to physicians who graduate from top-5 U.S. medical schools.[51] To understand how incomes shape specialty choice across the broader ability distribution, we use different data and a different empirical approach. Both methods yield similar results for the top (least-constrained) residency applicants, supporting our approach.

Score-Based Aggregate Model Specification. To account for the limited choices facing most residency applicants, we estimate an analogue of equation (14) allowing for heterogeneity by test score. To do this, we use aggregate data on the number of physicians who apply to each specialty grouped by USMLE score.[52] The National Residency Matching Program (NRMP) reports these scores in 10-point intervals, ranging from ≤ 180 to > 260. We simplify the model by removing other dimensions of heterogeneity from equation (14), which allows us

[50] Appendix C.3 has more details on these variables and estimation.

[51] Furthermore, unlike the canonical consumer demand model, our setting does not (directly) have strategically income-setting firms. Indeed, as Section 5 showed, much of the variation in income is driven by government fiat.

[52] This exam is normally taken after the second year of medical school and historically played an important role in the residency match. In 2022, numerical scores were replaced with pass/fail grading (USMLE, 2021).

to estimate the model on aggregate data using the log-shares transformation (Berry, 1994), treating family medicine as the outside option. We estimate a pooled specification using data for all USMLE score intervals, but allow the coefficient on mean hourly earnings in a specialty to vary across score intervals:

$$\ln \Pr(s)_{ast} - \ln \Pr(0)_{at} = \delta_a^M \cdot M_{st} + \delta^\phi \phi_{st} + \delta_s + \nu_{ast}. \tag{15}$$

In equation (15), $\ln \Pr(s)_{ast}$ denotes the log share of graduates in score group a in cohort t who applied to specialty s, and other variables are defined as in equation (14).

We use USMLE score buckets to proxy for ability a. If applications reflect rationing of attractive specialties, the empirical choice patterns should look different for physicians who are not at the top of the USMLE score distribution and hence more constrained. We expect δ_a^M for high USMLE scores to reflect true preferences, measuring how graduates trade off earnings and non-pecuniary amenities of specialties. As we move down in the USMLE score distribution, the coefficient on specialty earnings potential should shift, reflecting the shadow price of the entry constraint. At the bottom of the distribution, it could even reverse sign—not because lower-scoring graduates have different preferences, but because their choice is limited to the slots remaining after higher-scoring candidates make their selections.

Results. Table 5 reports estimates of equation (14)'s α parameter in the top row and each score group's δ_a^M from (15) in the bottom row.[53] The estimates for high-ability groups are generally intuitive: graduates who are likely unconstrained in their choices prefer higher earnings. Physicians prefer the amenities offered by primary care and various procedural specialties appear to have relative disamenities. Our estimate of $\hat{\alpha}$ in the highest USMLE score group based on the aggregate model is very similar to the individual choice model estimate. This lends credence to these numbers, as the two approaches use entirely different

[53] The δ_a^M reported in the bottom row are calculated as the sum of the group-specific interaction with hourly income, reported on the second row, plus the hourly income coefficient for the omitted group (score ≤ 190). Table E.11 reports detailed estimates for the individual and Table E.13 for the aggregate model.

datasets and estimation methods.

Moving across the columns to the right, we see the implied coefficient on income $\hat{\delta}_a^M$ for students with lower USMLE scores become smaller and even turn negative for the lowest score groups. This is consistent with the equilibrium described in our model in which higher compensation attracts new physicians to a specialty, while the scarcity of residency slots screens lower-ability physicians out of high-paying specialties. Appendix C.3 presents additional details and graphs to help interpret these results. The specialty choice elasticities are in a similar range to earlier studies that account for rationed entry into the highest-paid specialties (Nicholson, 2002; Bhattacharya, 2005).[54]

6.2 Implications for Allocation of Talent

We use these results to consider two salient policy debates surrounding physician labor markets: entry restrictions and shortages of primary care physicians (Glied et al., 2009). Entry restrictions may drive physician shortages, especially in rural areas, while student debt allegedly leads doctors to choose highly compensated specialties over primary care.

Policy discussions often consider increases to primary care physicians' incomes, either through reimbursements, bonuses, or loan forgiveness. We use our estimates to compute how physicians' specialty choice responds to these policies. Figure 8A shows the results for physicians who are likely unconstrained in their specialty choice, based on the individual-level model. Increasing primary care physicians' hourly income to the level of medicine subspecialists (an increase from \$98 to \$168) would induce 48% of top-5 graduates in the U.S. to practice primary care rather than specializing. This is a 20 percentage point increase. Nearly half of these reallocations would be from procedural specialties, followed by radiology and surgery. This policy would nearly double the share of primary care physicians coming from top-5 medical schools (Figure E.13). This implies that physicians value the amenities

[54]Consistent with studies that do not account for entry barriers and find much lower elasticities (Nicholson and Propper, 2011), we get much lower elasticity estimates as we move down the USMLE score distribution where MD graduates likely have less choice.

of primary care. About a third of today's specialists would have preferred primary care if its hourly income were the same as specialists'.[55]

Figure 8B reports a similar exercise using the aggregated model with heterogeneous test scores. We compute the counterfactual distribution of internists' test scores that results when internal medicine hourly earnings are increased to dermatologists' level. The share of graduates with USMLE scores above 250 in internal medicine increases by more than 10 percentage points, displacing some lower-scoring entrants, for a 10 point increase in the average USMLE score in internal medicine.

We frame these analyses as an increase in primary care physicians' incomes. Since the model considers relative incomes, the results would be the same if we instead reduced specialists' earnings.[56] This distinction—*i.e.* the absolute level of earnings—may affect the choice to enter medicine in the first place. While our model abstracts from this decision, Appendix D uses our data to speculate on implications for this extensive occupational choice margin.

7 Conclusion

This paper uses a new administrative data linkage to describe and understand U.S. physicians' earnings. Physicians' care commands one-fifth of healthcare resources, but their personal earnings account for less than half of that amount, totaling 8.6% of U.S. healthcare spending. Physicians earn $350,000 on average, for an average lifetime income of $10 million. The age-earnings profile is steep, reflecting the many years of human capital investments required to enter a career in medicine. We show that, in this setting with binding entry restrictions, government payment rules have a profound impact on earnings and thus play a key role in valuing and allocating one of society's most expensive assets: physicians' human

[55] Raising primary care income to the level of procedural specialists is in practice very similar to equalizing incomes across all specialities, see Figure E.14. Note that eliminating student loans, also frequently proposed as a way to induce new physicians to pursue primary care, is not equivalent to income equalization. As medical school debt varies little across specialties, it has little impact on relative incomes.

[56] Paying all physicians at the current level of U.S. primary care incomes would approximate physician compensation in Sweden, while paying everyone at U.S. specialists' level would be closer to Switzerland.

capital.

Our results teach how policy drives the most consequential long-run outcomes in this labor market and provide a clear agenda for future research. To analyze the long-run welfare impacts of healthcare policies, including those we investigate, we need evidence on the distribution of health impacts and thus social returns to physician ability in different specialties. We do not speculate on the magnitude of such returns in this paper, but our results show that quantifying the health impacts of ability is an essential direction for future work and is key to formulating optimal payment policies.

References

AAMC, "2022 FACTS: Applicants and Matriculants Data," 2022. Association of American Medical Colleges, Washington, D.C.

Abowd, John M., Francis Kramarz, and David N. Margolis, "High Wage Workers and High Wage Firms," *Econometrica*, 1999, *67* (2), 251–333.

Ahmed, Abdul-Kareem and Eli Y. Adashi, "Research Fever — An Ever More Prominent Trend in the Residency Match," *JAMA Surgery*, 01 2023, *158* (1), 6–8.

Altonji, Joseph and Ling Zhong, "The Labor Market Returns to Advanced Degrees," *Journal of Labor Economics*, 2021, *39* (2), 303–360.

Andrews, Martyn J., Len Gill, Thorsten Schank, and Richard Upward, "High Wage Workers and Low Wage Firms: Negative Assortative Matching or Limited Mobility Bias?," *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 2008, *171* (3), 673–697.

Bailey, Martha J., Hilary W. Hoynes, Maya Rossin-Slater, and Reed Walker, "Is the Social Safety Net a Long-Term Investment? Large-Scale Evidence from the Food Stamps Program," Working Paper No. 26942, National Bureau of Economic Research April 2020.

Baker, Dean, "The problem of doctors' salaries," *Politico*, October 25 2017. Available online at https://www.politico.com/agenda/story/2017/10/25/doctors-salaries-pay-disparities-000557/ (accessed June 28, 2023).

Baker, Laurence C., "Differences in Earnings between Male and Female Physicians," *New England Journal of Medicine*, 1996, *334* (15), 960–964.

Bell, Alex, Raj Chetty, Xavier Jaravel, Neviana Petkova, and John Van Reenen, "Do Tax Cuts Produce More Einsteins? The Impacts Of Financial Incentives Versus Exposure To Innovation On The Supply Of Inventors," *Journal of the European Economic Association*, 2019.

Berry, Steven T., "Estimating Discrete-Choice Models of Product Differentiation," *The RAND Journal of Economics*, 1994, pp. 242–262.

Bertrand, Marianne, "CEOs," *Annual Review of Economics*, 2009, *1* (1), 121–150.

Bhattacharya, Jayanta, "Specialty Selection and Lifetime Returns to Specialization within Medicine," *The Journal of Human Resources*, 2005, *40* (1), 115–143.

Bonhomme, Stéphane, Kerstin Holzheu, Thibaut Lamadon, Elena Manresa, Magne Mogstad, and Bradley Setzler, "How Much Should We Trust Estimates of Firm Effects and Worker Sorting?," *Journal of Labor Economics*, 2023, *41* (2).

Bresnahan, Timothy F. and Peter C. Reiss, "Entry and Competition in Concentrated Markets," *Journal of Political Economy*, 1991, *99* (5), 977–1009.

Cabral, Marika, Colleen Carey, and Sarah Miller, "The Impact of Provider Payments on Health Care Utilization of Low-Income Individuals: Evidence from Medicare and Medicaid," Working Paper No. 29471, National Bureau of Economic Research November 2021.

Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline, "Firms and Labor Market Inequality: Evidence and Some Theory," *Journal of Labor Economics*, January 2018, *36* (S1), 13–70.

—— **, Jesse Rothstein, and Moises Yi**, "Location, Location, Location," Working Paper CES 21-32, Center for Economic Studies October 2021. Available online at https://www2.census.gov/ces/wp/2021/CES-WP-21-32.pdf (accessed July 5, 2023).

Case, Anne and Angus Deaton, *Deaths of Despair and the Future of Capitalism*, Princeton: Princeton University Press, 2020.

CBO, "Federal Subsidies for Health Insurance Coverage for People Under Age 65: 2016 to 2026," March 2016. Congressional Budget Office. Available online at https://www.cbo.gov/sites/default/files/114th-congress-2015-2016/reports/51385-healthinsurancebaseline.pdf (accessed July 5, 2023).

Chan, David C. and Michael J. Dickstein, "Industry Input in Policy Making: Evidence from Medicare," *The Quarterly Journal of Economics*, 04 2019, *134* (3), 1299–1342.

Chen, M. Keith and Judith A. Chevalier, "Are Women Overinvesting in Education? Evidence from the Medical Profession," *Journal of Human Capital*, 2012, *6* (2), 124–149.

Chetty, Raj, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez, "Where is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States," *The Quarterly Journal of Economics*, 09 2014, *129* (4), 1553–1623.

Clemens, Jeffrey and Joshua D. Gottlieb, "Do Physicians' Financial Incentives Affect Treatment Patterns and Patient Health?," *American Economic Review*, April 2014, *104* (4), 1320–1349.

—— **and** ——, "In the Shadow of a Giant: Medicare's Influence on Private Payment Systems," *Journal of Political Economy*, February 2017, *125* (1), 1–39.

—— **,** —— **, and Tímea Laura Molnár**, "Do Health Insurers Innovate? Evidence From the Anatomy of Physician Payments," *Journal of Health Economics*, September 2017, *55*, 153–167.

CMS, "National Health Expenditure Tables," *Centers for Medicare and Medicaid Services*, 2019. Available online at https://www.cms.gov/Research-Statistics-Data-and-Systems/Statistics-Trends-and-Reports/NationalHealthExpendData/NationalHealthAccountsHistorical (accessed June 22, 2021).

Cooper, Zack, Stuart V. Craig, Martin Gaynor, and John Van Reenen, "The Price Ain't Right? Hospital Prices and Health Spending on the Privately Insured," *The Quarterly Journal of Economics*, 2019, *134* (1), 51–107.

Crawford, Jay, "My Hand Surgeon Should Have Been Paid $4.5 billion For Fixing My Broken Wrist, Not $1,000," *STAT News*, August 28 2019. Available online at https://www.statnews.com/2019/08/28/hand-surgeon-payment-entrepreneur-reward/ (accessed March 22, 2019).

Diamond, Rebecca and Enrico Moretti, "Where is Standard of Living the Highest? Local Prices and the Geography of Consumption," Working Paper No. 29533, National Bureau of Economic Research December 2021.

Dingel, Jonathan, Joshua D. Gottlieb, Maya Lozinski, and Pauline Mourot, "Market Size and Trade in Medical Services," Working Paper No. 31030, National Bureau of Economic Research March 2023.

Dunn, Abe and Adam Hale Shapiro, "Physician Market Power and Medical-Care Expenditures," *The Journal of Law and Economics*, 2014, *57* (1), 159–193.

Esteves-Sorenson, Constança and Jason Snyder, "The Gender Earnings Gap for Physicians and its Increase over Time," *Economics Letters*, 2012, *116* (1), 37–41.

Feldstein, Martin S., "The Rising Price of Physician's Services," *The Review of Economics and Statistics*, 1970, *52* (2), 121–133.

Finkelstein, Amy, "The Aggregate Effects of Health Insurance: Evidence from the Introduction of Medicare," *The Quarterly Journal of Economics*, 2007, *122* (1), 1–37.

—, Matthew Gentzkow, and Heidi Williams, "Sources of Geographic Variation in Health Care: Evidence from Patient Migration," *The Quarterly Journal of Economics*, 2016, *131* (4), 1681–1726.

—, —, and —, "Place-Based Drivers of Mortality: Evidence from Migration," *American Economic Review*, August 2021, *111* (8), 2697–2735.

Friedman, Milton and Simon Kuznets, *Income from Independent Professional Practice*, National Bureau of Economic Research, 1945.

Fuchs, Victor R. and Marcia J. Kramer, *Determinants of Expenditures for Physicians' Services in the United States 1948–68*, National Bureau of Economic Research, 1973.

Ganguli, Ishani, Bethany Sheridan, Joshua Gray, Michael Chernew, Meredith B Rosenthal, and Hannah Neprash, "Physician Work Hours and the Gender Pay Gap—Evidence from Primary Care," *New England Journal of Medicine*, 2020, *383* (14), 1349–1357.

GAO, "Medicare: Information on Geographic Adjustments to Physician Payments for Physicians' Time, Skills, and Effort," Report to Congressional Requesters GAO-22-103876, Government Accountability Office February 2022.

Garthwaite, Craig, Tal Gross, and Matthew Notowidigdo, "Hospitals as Insurers of Last Resort," *American Economic Journal: Applied Economics*, 2018, *10* (1), 1–39.

Glied, Sherry, Ashwin G. Prabhu, and Norman Edelman, "The Cost of Primary Care Doctors," *Forum for Health Economics & Policy*, 2009, *12* (1).

Gottlieb, Joshua D., David Hémous, Jeffrey Hicks, and Morten Olsen, "The Spillover Effects of Top Income Inequality," Working Paper No. 31366, National Bureau of Economic Research June 2023.

Hsieh, Chang-Tai, Erik Hurst, Charles I. Jones, and Peter J. Klenow, "The Allocation of Talent and US Economic Growth," *Econometrica*, 2019, *87* (5), 1439–1474.

Jagsi, Reshma, Kent A. Griffith, Abigail Stewart, Dana Sambuco, Rochelle DeCastro, and Peter A. Ubel, "Gender Differences in the Salaries of Physician Researchers," *JAMA*, June 2012, *307* (22), 2410–2417.

Jaimovich, Nir and Sergio Rebelo, "Nonlinear Effects of Taxation on Growth," *Journal of Political Economy*, 2017, *125* (1), 265–291.

Kline, Patrick, Neviana Petkova, Heidi Williams, and Owen Zidar, "Who Profits from Patents? Rent-Sharing at Innovative Firms," *The Quarterly Journal of Economics*, August 2019, *134* (3), 1343–1404.

—— **, Raffaele Saggio, and Mikkel Sølvsten**, "Leave-Out Estimation of Variance Components," *Econometrica*, 2020, *88* (5), 1859–1898.

LoSasso, Anthony T., David Armstrong, Gaetano Forte, and Susan E. Gerber, "Differences In Starting Pay For Male And Female Physicians Persist; Explanations For The Gender Gap Remain Elusive," *Health Affairs*, 2020, *39* (2), 256–263.

MaCurdy, Thomas, Jason Shafrin, Mary Beth Blasnek, Sajid Zaidi, Mallory Bounds, Brandy Lipton, David Pham, Sonam Sherpa, Elen Shrestha, Silvia Manolache et al., "Revised Final Report on the CY 2014 Update of the Geographic Practice Cost Index for the Medicare Physician Fee Schedule," 2014. Acumen, LLC. Available online at https://www.cms.gov/Medicare/Medicare-Fee-for-Service-Payment/PhysicianFeeSched/Downloads/CY2015-PFS-FR-GPCI.pdf (accessed July 5, 2023).

Molitor, David, "The Evolution of Physician Practice Styles: Evidence From Cardiologist Migration," *American Economic Journal: Economic Policy*, 2018, *10* (1), 326–56.

Murphy, Kevin M., Andrei Shleifer, and Robert W. Vishny, "The Allocation of Talent: Implications for Growth," *The Quarterly Journal of Economics*, 1991, *106* (2), 503–530.

Nicholson, Sean, "Physician Specialty Choice under Uncertainty," *Journal of Labor Economics*, 2002, *20* (4), 816–847.

—— **and Carol Propper**, "Medical Workforce," in Mark V. Pauly, Thomas G. Mcguire, and Pedro P. Barros, eds., *Handbook of Health Economics*, Vol. 2, Elsevier, 2011, chapter 14, pp. 873–925.

—— **and Nicholas S. Souleles**, "Physician Income Expectations and Specialty Choice," Working Paper No. 8536, National Bureau of Economic Research October 2001.

NRMP, "Data Release and Research Committee: Results of the 2022 NRMP Program Director Survey," *National Resident Matching Program, Washington, D.C.*, 2022.

Polyakova, Maria, Victoria Udalova, Geoffrey Kocks, Katie Genadek, Keith Finlay, and Amy N. Finkelstein, "Racial Disparities In Excess All-Cause Mortality During The Early COVID-19 Pandemic Varied Substantially Across States," *Health Affairs*, 2021, *40* (2), 307–316.

Seabury, Seth A., Amitabh Chandra, and Anupam B. Jena, "Trends in the Earnings of Male and Female Health Care Professionals in the United States, 1987 to 2010," *JAMA Internal Medicine*, 10 2013, *173* (18), 1748–1750.

Sloan, Frank A., "Physician Supply Behavior in the Short Run," *ILR Review*, 1975, *28* (4), 549–569.

Smith, Matthew, Danny Yagan, Owen Zidar, and Eric Zwick, "Capitalists in the Twenty-First Century," *The Quarterly Journal of Economics*, July 2019, *134* (4), 1675–1745.

Tolbert, Jennifer, Kendal Orgera, and Damico Anthony, "Key Facts about the Uninsured Population," *Kaiser Family Foundation*, 2020. Available online at https://www.kff.org/uninsured/issue-brief/key-facts-about-the-uninsured-population/ (accessed June 18, 2021).

USMLE, "USMLE Step 1 Transition to Pass/Fail Only Score Reporting," September 15 2021. Available online at https://web.archive.org/web/20230525135028/https://www.usmle.org/usmle-step-1-transition-passfail-only-score-reporting (accessed September 15, 2021).

Vaughn, Bryan T., Steven R. DeVrieze, Shelby D. Reed, and Kevin A. Schulman, "Can We Close The Income And Wealth Gap Between Specialists And Primary Care Physicians?," *Health Affairs*, 2010, *29* (5), 933–940.

Figure 1: **Physician Earnings over the Lifecycle and by Firm Size**

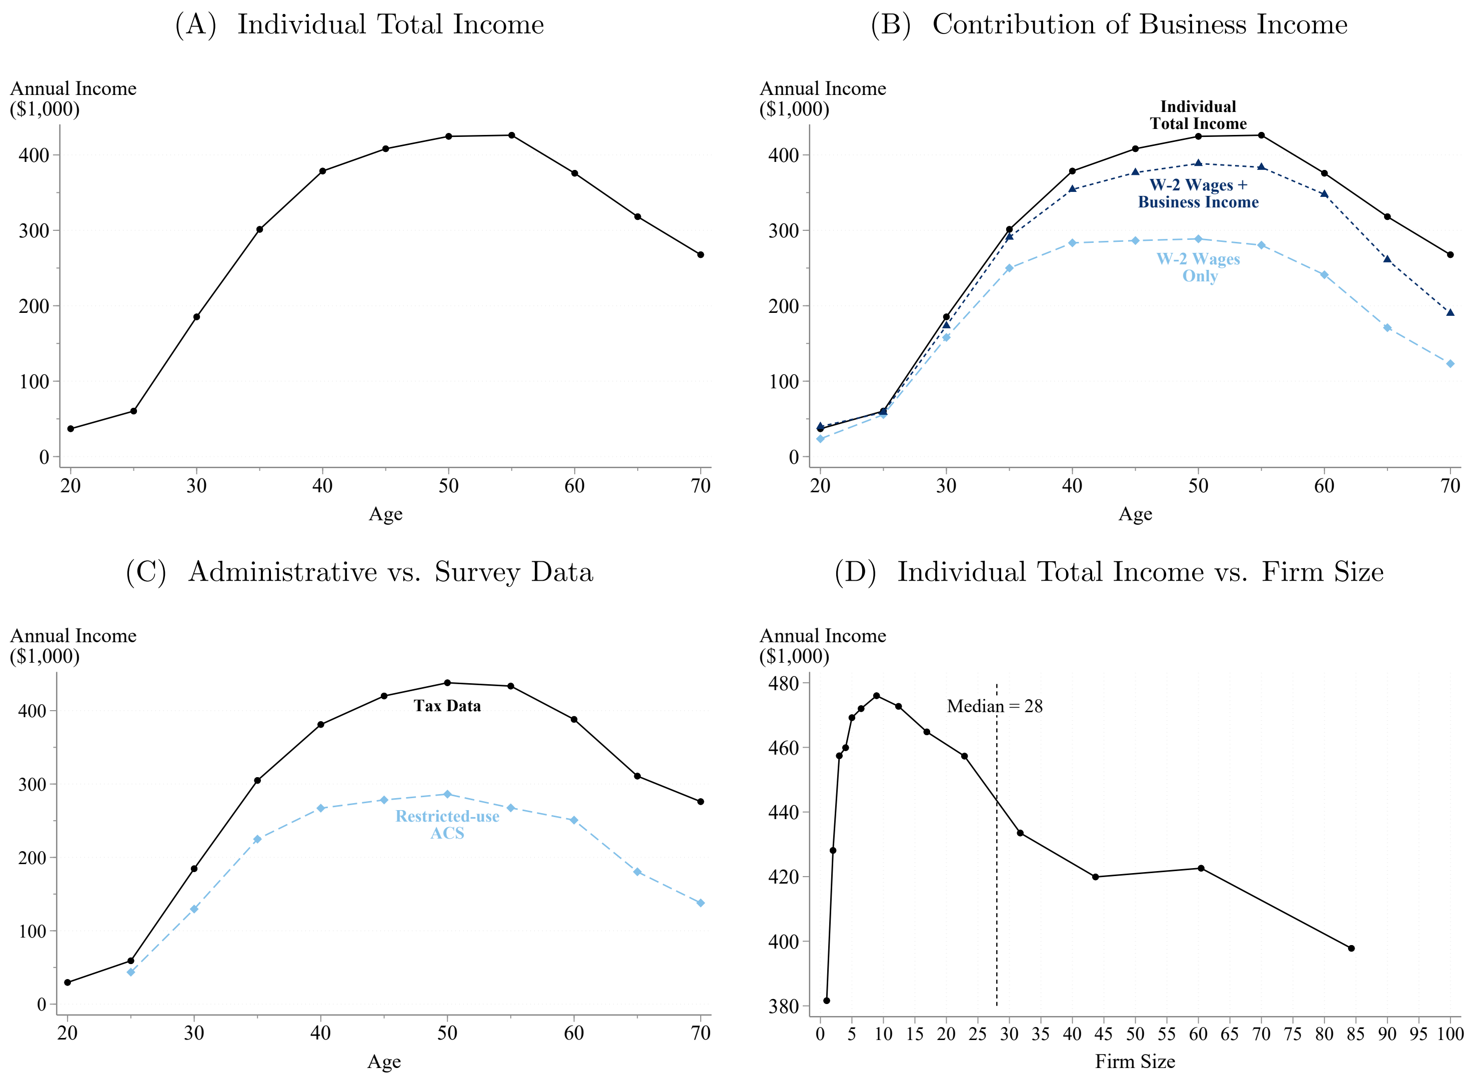

Notes: The figure plots mean individual total income in our 2017 sample of physicians by 5-year age intervals (Panels A-C) and by firm size (Panel D). Business income in Panel B is defined as the Total Money Income of the household net of wages, taxable dividends, taxable interest, social security, partially observed profit and loss from Schedule E, and distributions from pre-tax deferral accounts irrespective of age. ACS total individual income in Panel C is defined as the sum of individual wage and self-employment income of the index individual plus self-employment income of the spouse. Panel D is restricted to physicians age 40 to 55 and firms with fewer than 100 physicians; the x-axis is ventiles of the physician-level distribution of firm size. The term “firm” refers to the tax unit, measured as the EIN on Form W-2. Appendix B.2 provides measurement details. Disclosure Review Board approval CBDRB-FY23-0319, CBDRB-FY2023-CES005-024.

Figure 2: **Correlates of Specialty Income**

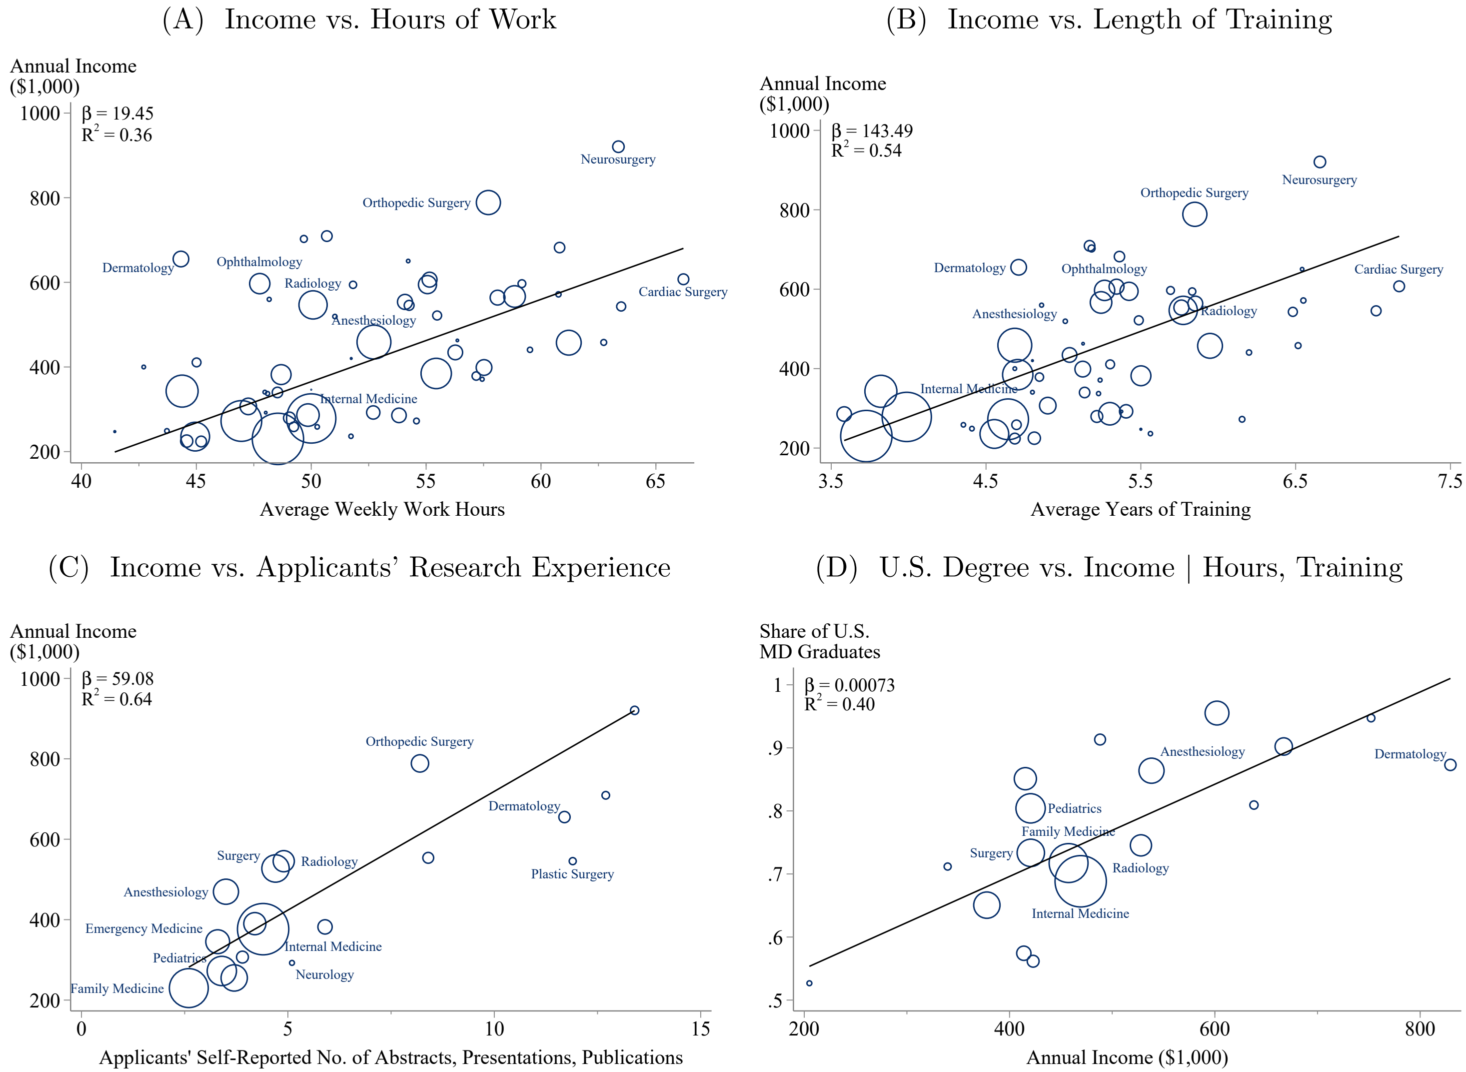

Notes: This figure characterizes the relationship between specialty earnings and specialty characteristics. Specialty earnings are measured as mean individual total income among 40- to 55-year-old physicians in our full panel 2005-2017 in a Medicare Specialty (Panels A and B) or NRMP specialty (Panels C and D). We plot specialty earnings against the average number of hours worked among physicians aged 40 to 55 in 2005–2017 (Panel A), the average imputed years of training (Panel B), and the average number of abstracts, presentations, and publications that MD students report having completed on their residency application, as provided by NRMP (Panel C). Panel D plots the specialty's share of physicians with a U.S. degree against average earnings, conditional on the number of work hours and years of training. Years of training is imputed from tax data as described in Appendix B.2. Circle sizes in the graphs are proportional to the number of individuals in each specialty in our baseline sample in 2017. The line of best fit is estimated as a weighted bivariate OLS on specialty-level data. Disclosure Review Board approval CBDRB-FY23-0319, CBDRB-FY2023-CES005-024.

Figure 3: **Geographic Variation in Earnings**

(A) Physicians

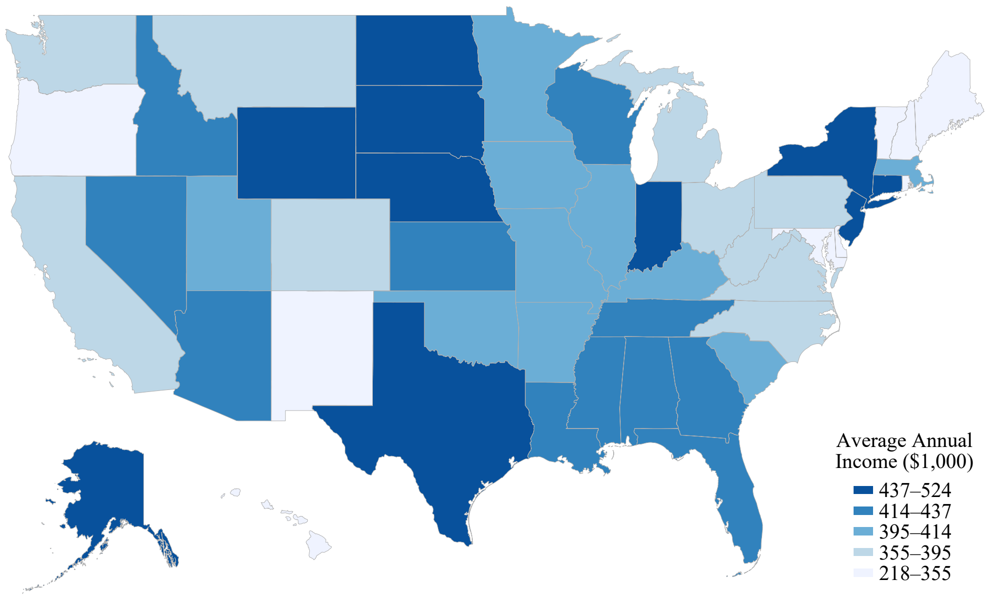

(B) Lawyers

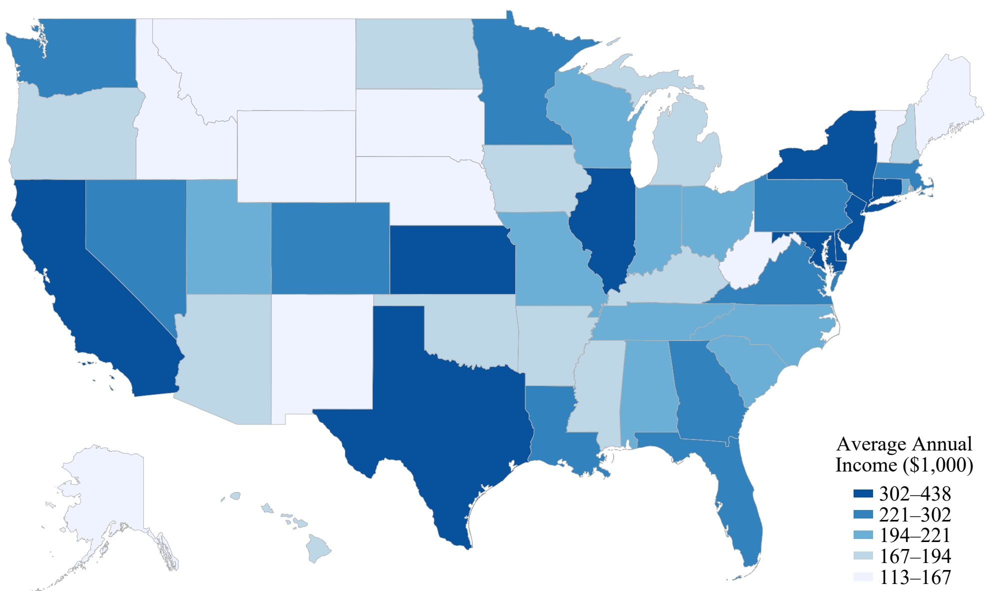

Notes: This figure plots mean individual total income among 40 to 55 year old physicians (Panel A) and lawyers (Panel B) in year 2017 by state. Income is measured using individual tax returns data and is defined as the sum of individual total wage income and the household AGI net of all wage earnings and taxable retirement distributions (for those aged 60 or older), but gross of tax exempt interest and social security payments. Physicians and lawyers are defined as described in Section 2 and Appendix B.2. The same Appendix also provides more details on income measurement. Disclosure Review Board approval CBDRB-FY23-0319, CBDRB-FY2023-CES005-024.

Figure 4: **Event Study: Physician Movers**

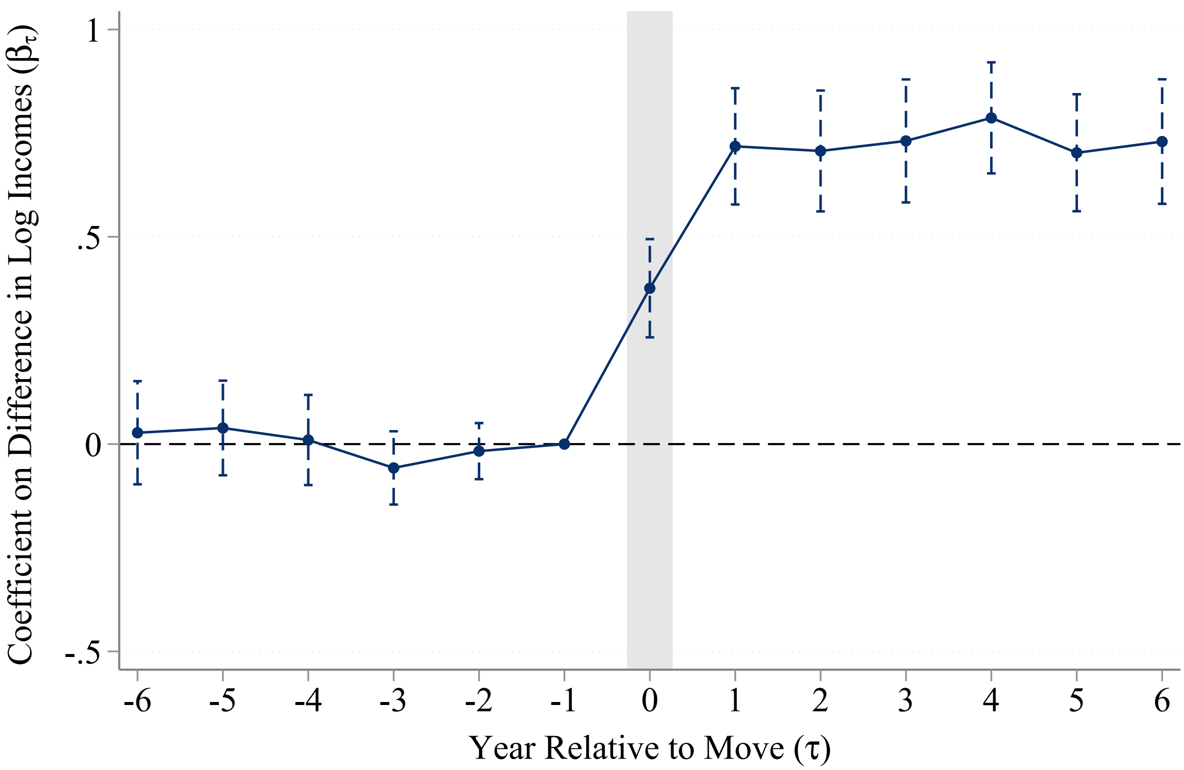

Notes: This figure shows coefficient estimates on the difference between mean individual total income between origin and destination commuting zones ($\Delta\overline{\ln y}_{(j,j')}$) from equation (7). The coefficient is normalized to 0 in the year prior to the move ($\tau = -1$). The dashed lines mark the 95% confidence intervals. The outcome variable is log individual total income. The independent variables include $\Delta\overline{\ln y}_{(j,j')}$ interacted with year effects, physician fixed effects, relative year fixed effects, and age fixed effects. A physician is defined as a mover and is included in the sample if they changed their commuting zone once between years 2005 to 2017, and were aged 40 to 55 during that change. Disclosure Review Board approval CBDRB-FY23-0319, CBDRB-FY2023-CES005-024.

Figure 5: **Place vs. Physician Contributions to Earnings**

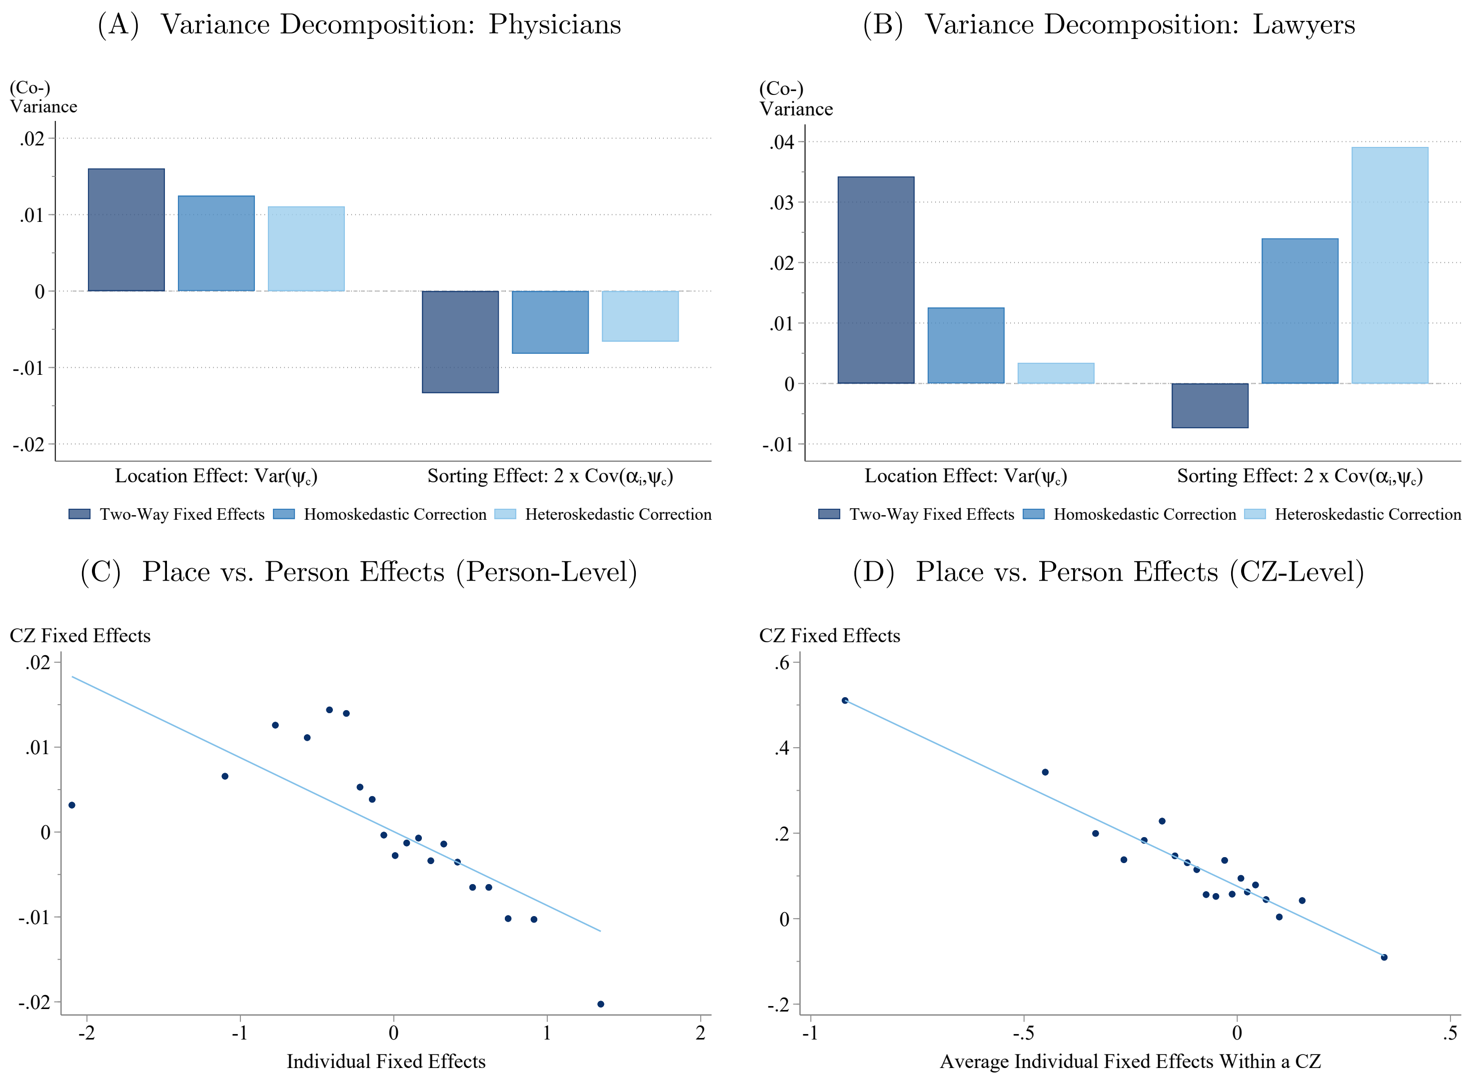

Notes: This figure shows elements of variance decomposition of individual total income among 40–55 year old physicians (Panel A) and lawyers (Panel B) in the sample of movers (see definition in Figure 4.) Estimates in bars labeled "Two-Way Fixed Effects" are based on the estimation of equation (8). The outcome variable is log individual total income. The importance of location effects is computed as the variance of estimated CZ fixed effects, $\text{Var}(\psi_c)$. The effect of sorting of people to locations, $2\text{Cov}(\alpha_i, \psi_c)$, is computed as twice the covariance of individual and CZ fixed effect estimates. The bars labeled homoskedastic and heteroskedastic correction report the corrected variance and covariance terms based on Andrews et al. (2008) and Kline et al. (2020), respectively, implemented following Bonhomme et al. (2023). Panels C and D show binned scatterplots relating place effects and person effects based on estimation of equation (8) in the sample of movers. Panel C reports the average CZ fixed effect within each ventile of individual fixed effects distribution. In Panel D we first collapse the data to the CZ level by averaging individual fixed effects within a CZ as in Card et al. (2021). The line of best fit is based on a bivariate OLS regression using underlying data points. Disclosure Review Board approval CBDRB-FY23-0319, CBDRB-FY2023-CES005-024.

Figure 6: **Correlates of Place Effects**

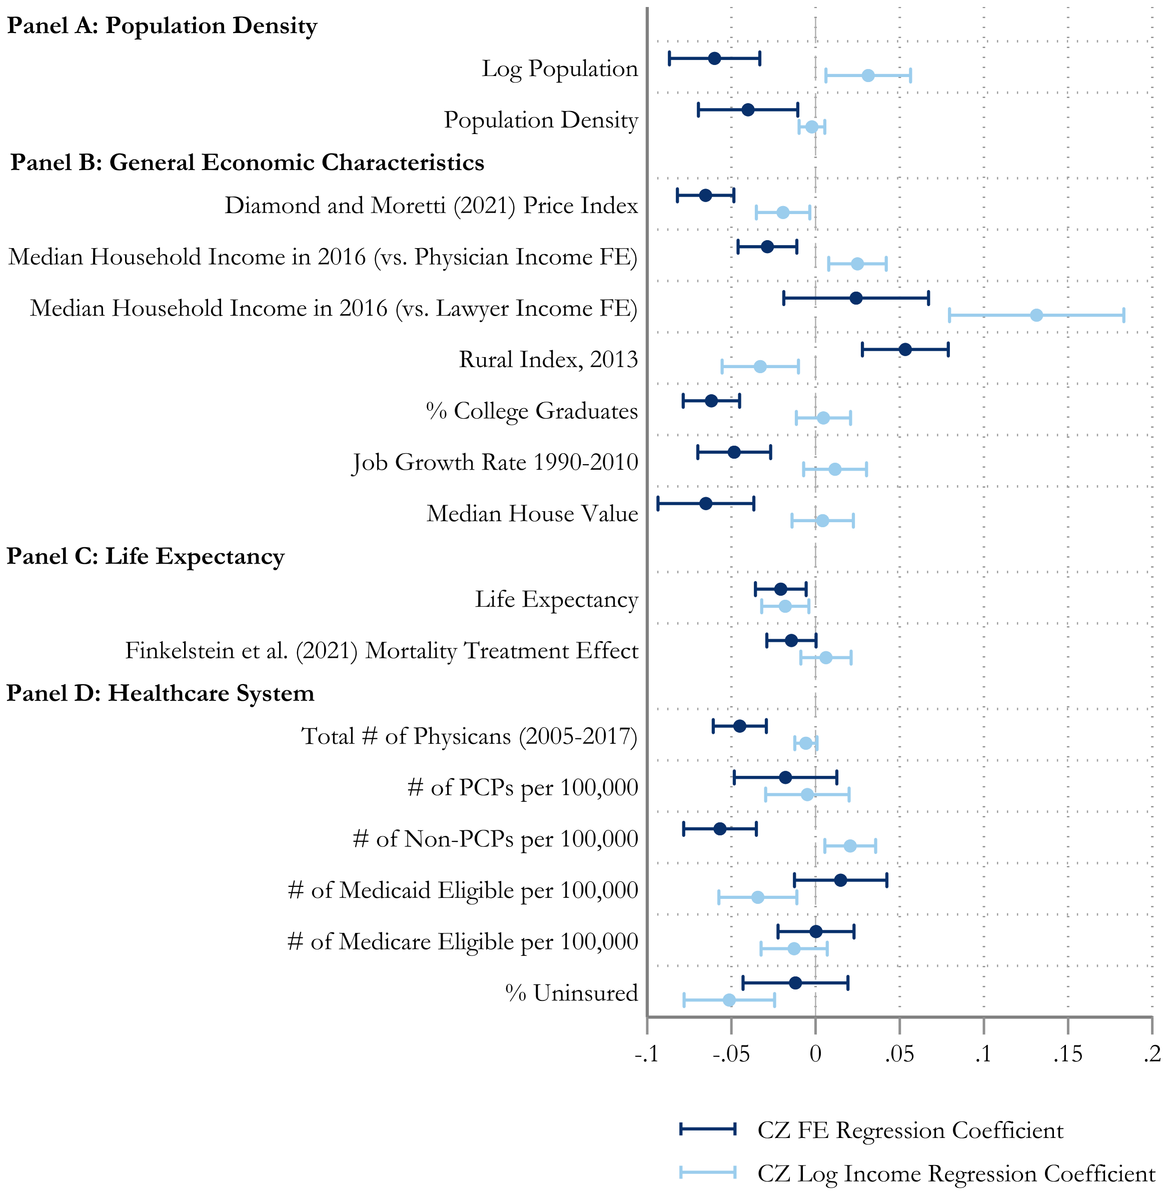

Notes: This figure plots the results of bivariate OLS regressions of raw average individual total income in a commuting zone (light blue colors), as well as of place treatment effect on earnings (dark blue), on z-scores of the indicated place characteristics. Place treatment effects on earnings are CZ fixed effects from the estimation of equation (8) in the sample of movers (see definition in Figure 4.). Raw mean income is computed in the same sample. CZ-level characteristics are as reported in Chetty et al. (2014); Finkelstein et al. (2021); Diamond and Moretti (2021). Disclosure Review Board approval CBDRB-FY23-0319, CBDRB-FY2023-CES005-024.

Figure 7: **Effect of ACA Insurance Expansion**

(A) First Stage: Share Insured

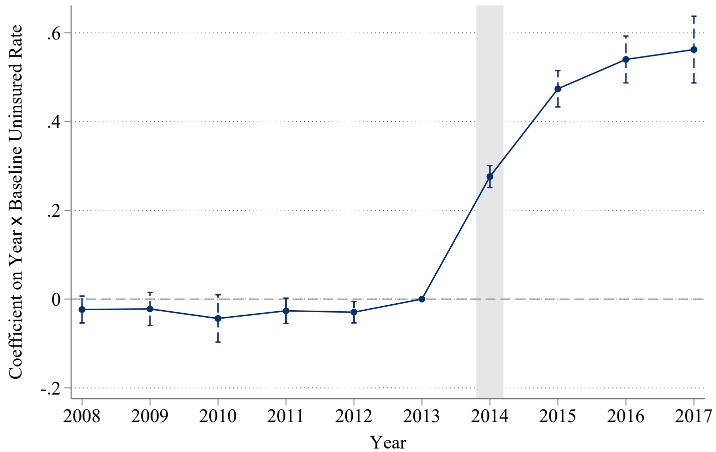

(B) Individual Total Income

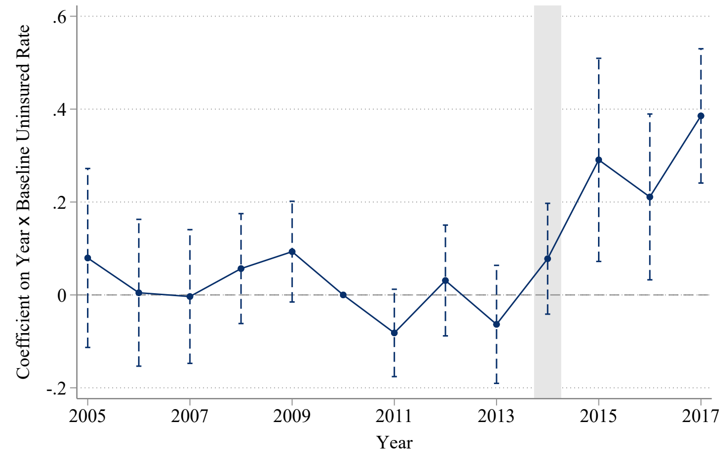

(C) Retirement

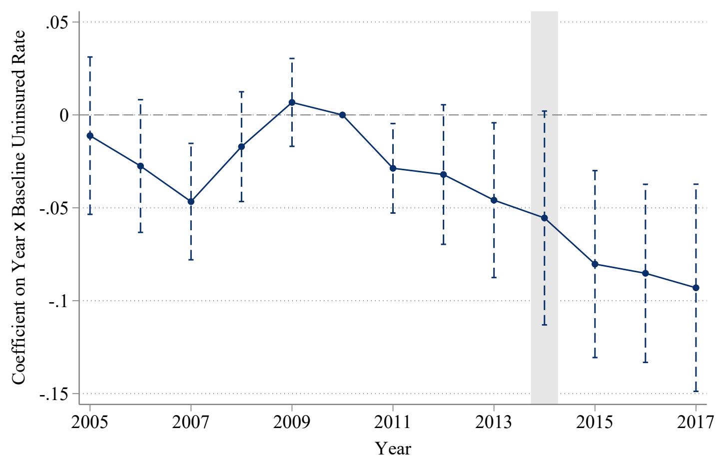

Notes: This figure shows event study estimates of the effects of ACA insurance expansions on insurance rates (Panel A), log individual total income of physicians aged 40–55 (Panel B), and the probability of retirement (defined as receiving Form 1099-SSA) among 56–70 year old physicians (Panel C). Independent variables include county fixed effects, year fixed effects, share of physicians of each age, and year fixed effects interacted with the share of under 65 population that was uninsured in 2013. The sample includes counties in states that had ACA expansions in 2014 and 2015 as detailed in Appendix C.2. The regression specification is in equation (13). Each regression is estimated on county-level data, weighted by the number of physicians (overall in Panel A and in the corresponding age group in Panels B and C) in that county and year. Error bars represent 95% confidence intervals. Disclosure Review Board approval CBDRB-FY23-0319, CBDRB-FY2023-CES005-024.

Figure 8: **Counterfactual Specialty Choices**

(A) Increase Primary Care Income to Medicine Subspecialty Level

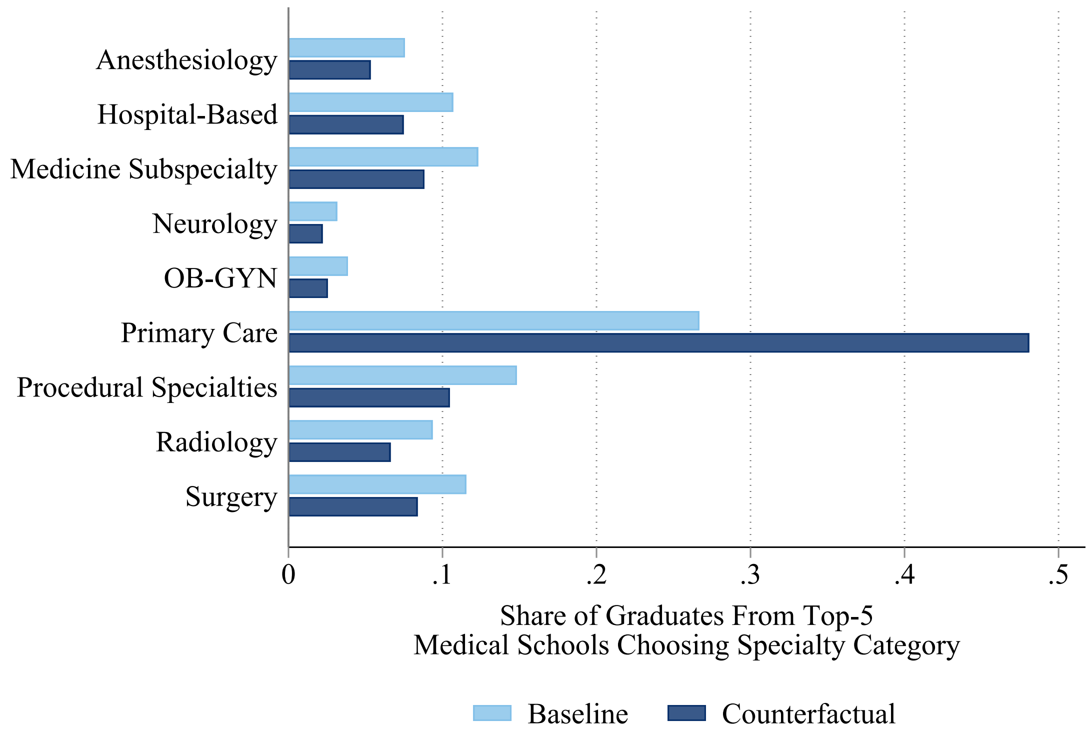

(B) Increase Internal Medicine Income to Dermatology Level

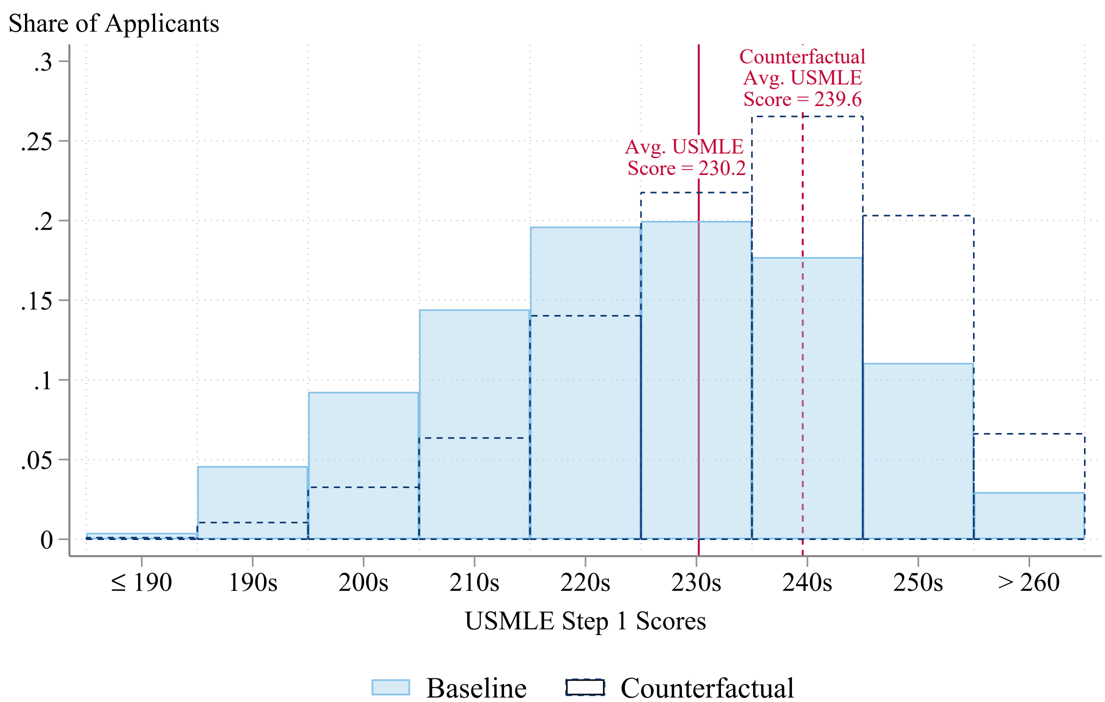

Notes: Panel A reports the results of a counterfactual exercise in which we set the mean hourly income of primary care physicians to be the same as that of medicine subspecialists. This counterfactual simulation uses the individual level discrete choice model model in equation (14) of Section 6.1. Mean hourly income is constructed as described in Appendix B.2. We plot the observed baseline and counterfactual share of graduates from top-5 MD programs who choose each specialty category. Panel B reports the results of a counterfactual in which we set the mean hourly income in internal medicine to equal the mean hourly income in dermatology. This computation uses the aggregate model in equation (15). We plot the observed and counterfactual distribution of USMLE scores in internal medicine. In both panels, the baseline is computed as model predictions at observed income values. Disclosure Review Board approval CBDRB-FY23-0319, CBDRB-FY2023-CES005-024.

Table 1: **Summary Statistics**

	Years:	2005-2017	2017		
		(1)	(2)	(3)	(4)
	Ages:	All	All	40 to 55	56 to 70
Number of Person-Years		11,600,000	848,000	350,000	287,000
Number of Unique Individuals		965,000	848,000	350,000	287,000
Demographics					
Age	Mean	45.3	49.3	47.3	62.6
	Median	45.0	49.0	47.0	62.0
	Std. Dev.	12.7	11.6	4.6	4.2
Female		0.36	0.38	0.40	0.26
Non-U.S.-Born		0.19	0.22	0.27	0.12
Married		0.77	0.80	0.83	0.82
Share Observed in ACS		0.20	0.20	0.20	0.22
Income					
Individual Total Wage (2017 $)	Mean	201,600	243,400	286,200	224,900
	Median	155,700	209,400	247,700	188,700
	Std. Dev.	945,300	283,000	260,300	359,400
Individual Total Income (2017 $)	Mean	290,800	350,000	404,500	367,500
	Median	210,700	265,000	308,600	267,500
	Std. Dev.	3,589,000	1,192,000	711,400	1,578,000
AGI (2017 $)	Mean	359,200	429,500	502,400	435,100
	Median	264,700	325,500	384,300	314,800
	Std. Dev.	859,200	1,266,000	827,500	1,652,000
Business Income > $25K		0.29	0.32	0.35	0.38
Households in Top 1% of AGI		0.22	0.24	0.31	0.24
Career Choices and Characteristics					
Firm Size (Number of Physicians)	Mean	1,101	1,472	1,536	1,480
	Median	52.0	84.6	75.0	20.0
	Std. Dev.	3,677	4,855	5,025	5,312
Weekly Working Hours (ACS)	Mean	50.5	49.5	49.5	47.6
	Median	50.0	50.0	50.0	50.0
	Std. Dev.	16.4	15.2	14.5	15.4
Retired (Based on 1099-SSA)		0.05	0.07	0.01	0.19
Name of Medical School Observed		0.52	0.53	0.56	0.55
Graduated from Ranked Medical School		0.48	0.48	0.48	0.50
Graduated from Top-5 Medical School		0.06	0.06	0.06	0.06
Share in Specialty Category					
Primary Care		0.43	0.44	0.42	0.43
Medicine Subspecialty		0.12	0.12	0.13	0.13
Hospital-Based		0.11	0.11	0.11	0.09
Surgery		0.09	0.09	0.09	0.09
Procedural Specialties		0.06	0.06	0.06	0.06
Anesthesiology		0.06	0.06	0.06	0.06
Radiology		0.05	0.05	0.05	0.05
OB-GYN		0.05	0.05	0.05	0.05
Neurology		0.03	0.03	0.03	0.03

Notes: This table reports summary statistics for all main samples used in our analysis. Column (1) includes years 2005–2017 and physicians aged 20 to 70. Columns (2) to (4) report summary statistics for the 2017 cross-section, overall, and by age subgroups. The sample in column (1) is constructed by merging the 2017 vintage of the National Plan and Provider Enumeration System (NPPES) file that includes National Provider Identifiers of all physicians in the U.S. with the universe of individual income tax return data. Section 2 and Appendix B.2 provide details on data sources and measurement of each variable. Disclosure Review Board approval CBDRB-FY23-0319, CBDRB-FY2023-CES005-024.

Table 2: **Characteristics of Top Earning Physicians**

		Top X% of Physicians by Income					
		1%	5%	10%	25%	50%	All
Number of Unique Individuals		3,500	17,500	35,000	87,500	175,000	350,000
Income and Labor Supply							
Individual Total Income ($1,000)	Mean	4,051	1,817	1,319	871	626	405
	Median	2,739	1,280	960	652	473	309
	Cutoff	1,937	960	719	473	309	-
Wage Income ($1,000)	Mean	897	737	654	524	417	286
AGI ($1,000)	Mean	4,465	1,993	1,448	964	708	502
Business Income ($1,000)	Mean	1,313	588	394	228	147	87
	Share > $25K	0.80	0.72	0.65	0.53	0.44	0.35
Median Share of Income from Business		0.28	0.20	0.13	0.05	0.02	0.00
Median Share of Income from Non-Labor		0.85	0.51	0.31	0.14	0.08	0.06
Median Share of Income from Labor		0.15	0.49	0.69	0.86	0.92	0.94
Mean Weekly Hours Worked		48	54	54	54	53	50
Retired (Based on 1099-SSA)		0.002	0.001	0.001	0.001	0.001	0.004
Mean Firm Size		354	449	493	699	1,091	1,536
Specialties and MD Training							
Graduated from Top-5 MD Program		0.10	0.08	0.07	0.07	0.06	0.06
Cardiology Share		0.02	0.04	0.05	0.05	0.03	0.02
Neurosurgery Share		0.06	0.05	0.04	0.02	0.01	0.01
General Surgery Share		0.03	0.03	0.04	0.05	0.05	0.03
Primary Care Share		0.16	0.12	0.12	0.15	0.24	0.42
Family Practice Share		0.03	0.03	0.03	0.04	0.07	0.14
Demographics							
Mean Age		48	48	48	48	48	47
Female		0.24	0.18	0.18	0.20	0.27	0.40
Non-U.S.-Born		0.22	0.23	0.23	0.23	0.25	0.27
Married		0.92	0.91	0.91	0.89	0.87	0.83
Share in New York and New Jersey		0.20	0.14	0.12	0.10	0.10	0.10
Share in California		0.11	0.09	0.09	0.10	0.12	0.12
Share in Florida		0.10	0.08	0.08	0.07	0.06	0.07
Share in Texas		0.11	0.11	0.10	0.09	0.08	0.08
Share in Arizona		0.03	0.03	0.02	0.02	0.02	0.02

Notes: This table reports selected summary statistics for the sample of age 40–55 physicians in year 2017 (sample in column (3) of Table 1), by selected percentiles of the individual total income distribution (as specified in column titles). Variables are as defined in Table 1. Section 2 and Appendix B.2 provide more details on data sources and measurement of each variable. Disclosure Review Board approval CBDRB-FY23-0319, CBDRB-FY2023-CES005-024.

Table 3: **RVU Regression Table**

	NPI-Level			Procedure-level		2SLS	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dependent variable:	Log Income	Log Total RVUs Billed	Log Number of Unique Procedures	Log Number of Unique Patients	Log Total RVUs Billed	Log Income	Retired
Panel A: Physicians Age 40-55							
Log Medicare Price Instrument ($\ln P_{i,t}$)	0.236	1.437	0.395				
	(0.035)	(0.109)	(0.039)				
Log RVUs per Procedure ($\ln RVU_{k,t}$)				0.344	1.382		
				(0.050)	(0.057)		
Log Total RVUs Billed ($\ln Q_{i,t}$)						0.167	
						(0.028)	
Log Income							-0.001
							(0.003)
Mean of Dependent Variable (2010-13)	13.13	8.75	2.99	3.88	4.82	13.13	0.00
Std. Dev. of Dependent Variable (2010-13)	0.84	1.13	0.79	1.12	1.70	0.84	0.06
Mean of Independent Variable	8.94	8.95	8.94	0.63	0.63	8.75	13.13
Std. Dev. of Independent Variable	1.02	1.00	1.02	1.18	1.18	1.12	0.84
Number of Observations	1,357,000	1,354,000	1,373,000	16,900,000	16,900,000	1,338,000	1,357,000
Panel B: Physicians Age 56-70							
Log Medicare Price Instrument ($\ln P_{i,t}$)	0.340	1.402	0.370				
	(0.042)	(0.107)	(0.044)				
Log RVUs per Procedure ($\ln RVU_{k,t}$)				0.394	1.441		
				(0.058)	(0.070)		
Log Total RVUs Billed ($\ln Q_{i,t}$)						0.246	
						(0.023)	
Log Income							-0.054
							(0.028)
Mean of Dependent Variable (2010-13)	13.02	8.68	2.96	3.90	4.78	13.02	0.10
Std. Dev. of Dependent Variable (2010-13)	0.88	1.07	0.78	1.13	1.69	0.88	0.30
Mean of Independent Variable	8.86	8.87	8.86	0.54	0.54	8.68	13.02
Std. Dev. of Independent Variable	0.96	0.94	0.96	1.13	1.13	1.06	0.88
Number of Observations	897,000	907,000	920,000	10,800,000	10,800,000	884,000	897,000

Notes: This table reports coefficients and standard errors from estimating equation (11) for each outcome variable as indicated in column names, and each age group, as indicated in panel names. Independent variables are the log Relative Value Units (RVU) rate, age fixed effects, and Medicare specialty by year fixed effects. For physician-level regressions, the log Medicare price ($\ln P_{i,j}$) faced by the physician is computed as a weighted average of procedure-level RVU rates for a fixed vector of services. 2SLS specifications regress the outcome variable of interest on the log total number of RVUs billed (RVU rate for each service multiplied by the number of times a service is performed) instrumented by $\ln P_{i,j}$. This is defined in equation (10), with the fixed vector of services defined as the average number of times each service (a combination of HCPCS procedure code and facility or non-facility place of service) was performed by a physician between years 2012 and 2017. Disclosure Review Board approval CBDRB-FY23-0319, CBDRB-FY2023-CES005-024.

Table 4: **ACA Regression Table**

	First Stage	Reduced Form			2SLS		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dependent variable:	Share Insured	Log Income	Share with Schedule SE	Share Retired	Log Income	Share with Schedule SE	Share Retired
Share Uninsured in 2013 ($U_{c,2013}$)							
$\times$ Years $2010-2013$	0.013	-0.076	0.082	-0.021			
	(0.011)	(0.046)	(0.037)	(0.019)			
$\times$ Year ≥ 2014	0.496	0.204	0.199	-0.064			
	(0.039)	(0.064)	(0.044)	(0.027)			
Share Insured ($I_{c,t}$)					0.487	0.323	-0.107
					(0.119)	(0.067)	(0.037)
Mean of Dependent Variable	0.851	12.470	0.469	0.199	12.470	0.456	0.198
Std. Dev. of Dependent Variable	0.047	0.158	0.078	0.051	0.154	0.078	0.049
Mean of Independent Variable	0.147	0.147	0.147	0.150	0.876	0.876	0.877
Std. Dev. of Independent Variable	0.044	0.044	0.044	0.045	0.051	0.051	0.052
Number of Observations	11,500	15,000	15,000	15,000	11,500	11,500	11,500
Number of Unique Counties	1,200	1,200	1,200	1,200	1,200	1,200	1,200
Physician Age Range	40-55	40-55	40-55	56-70	40-55	40-55	56-70

Notes: The table displays parametric difference-in-differences estimates of the effects of the ACA insurance expansions on the outcomes indicated in column names. The regression specification is as in equation (13), except that we collapse the time dimension into three periods: before the ACA passage (2010 and earlier); post-ACA passage and pre-implementation period (2011-2013); and post-implementation period (2014-2017). Age range restrictions are specified in the last row of the table. Independent variables include these three time intervals interacted with the fraction of population that was uninsured in a county in 2013, as well as county fixed effects, calendar year fixed effects, and the time-varying share of physicians of each age in each county. All regressions are estimated at the county-year level with the number of physicians in each county-year in the corresponding sample as regression weights. Standard errors are clustered at the county level. County-level averages are computed for physicians in our baseline sample who resided in states that expanded Medicaid in 2014 or 2015. Appendix C.2 provides the full list of states. Column (1) reports the first stage, where the outcome variable is the share of individuals under 65 who are insured in a county. Columns (2) to (4) report reduced-form estimates. Columns (5) to (7) report the results of corresponding 2SLS specifications that treat the rate of insurance in the under-65 population as the endogenous variable of interest and the rate of uninsured population in 2013 as an instrument. The 2SLS specification treats all years pre-implementation as the pre-period. Table E.9 reports the same specifications, but dropping the post-ACA passage and pre-implementation years (2011-2013). Disclosure Review Board approval CBDRB-FY23-0319, CBDRB-FY2023-CES005-024.

Table 5: **Specialty Choice Model**

Ability Group (a) :	Top-5 Ranked Medical School	USLME Step 1 Scores								
		> 260	251-260	241-250	231-240	221-230	211-220	201-210	191-200	≤ 190
Coefficient on Hourly Income (α)	0.014 (0.003)									
Coefficient on Hourly Income $\times$ Ability Group Dummy		0.025 (0.001)	0.022 (0.001)	0.019 (0.001)	0.014 (0.001)	0.010 (0.001)	0.005 (0.001)	0.003 (0.001)	-0.001 (0.001)	Reference -
Implied Coefficient on Hourly Income for Group a (δ_a^M)		0.016	0.014	0.011	0.006	0.002	-0.003	-0.005	-0.009	-0.008 (0.002)

Notes: This table reports selected coefficients from estimating the discrete choice model specified in Section 6.1. The estimates in column (1) are based on the model in equation (14), estimated on individual-level data using simulated maximum likelihood. The estimates in columns (2) to (10) are based on the group level model in equation (15) and are estimated on group data at the USMLE Step 1 score group-by-specialty level using OLS. Column (1) reports the estimate of $\hat{\alpha}$ for graduates from top-5 medical schools. The coefficients in columns (2) to (10) are estimates from one pooled regression. We first report the estimated interactions between hourly income and score group dummies and compute δ_a^M by adding interactions with the reference group estimate to get a score-group specific marginal utility of income coefficient. See Section 6.1 for more discussion of the interpretation. Table E.11 reports the full set of estimates for the specification in column (1). Table E.13 reports the full set of estimates for the specification in columns (2) to (10). Table E.12 presents own and cross-income elasticities of specialty choice probability computed based on the individual-level model of column (1). The elasticities are computed from model simulations. Table E.14 reports the analogue based on the aggregate model specification. Disclosure Review Board approval CBDRB-FY23-0319, CBDRB-FY2023-CES005-024.