Submitted to *Econometrica*

THE SPILLOVER EFFECTS OF TOP INCOME INEQUALITY

JOSHUA D. GOTTLIEB
Harris School of Public Policy, University of Chicago and NBER

DAVID HÉMOUS
Department of Economics, University of Zürich and CEPR

JEFFREY HICKS
Department of Management, University of Toronto

MORTEN OLSEN
Department of Economics, University of Copenhagen

Top income inequality in the United States has increased considerably within many occupations. This phenomenon has led to a search for a common explanation. We instead develop a theory where increases in income inequality originating within a few occupations can "spill over" through consumption into others. We show theoretically that such spillovers occur when an occupation provides non-divisible services of heterogeneous quality to consumers. Examining local income inequality across U.S. regions, we find evidence that such spillovers exist for physicians, dentists, other medical occupations, and real estate agents. Estimated spillovers for other occupations are consistent with the predictions of our theory. Calibrating our model, we show that spillovers amplify a given shock to top income inequality by at least 16%. Spillovers dampen the increase in consumers' welfare inequality compared with the change in income inequality.

JEL CLASSIFICATION. D31, J24, J31, O15.
KEYWORDS: Income inequality, Assignment model, Occupational inequality, Superstars.

1. INTRODUCTION

The increase in top earnings since the 1980s has been accompanied by growing inequality within the top of the distribution, both in aggregate and within occupations (Bakija, Cole,

Joshua D. Gottlieb: jgottlieb@uchicago.edu
David Hémous: david.hemous@econ.uzh.ch
Jeffrey Hicks: jeffrey.hicks@utoronto.ca
Morten Olsen: mgo@econ.ku.dk

We are deeply indebted to Jeffrey Clemens for his major role collaborating with us on earlier versions of this paper. Olsen gratefully acknowledges the financial support of the European Commission under the Marie Curie Research Fellowship program (Grant Agreement PCIG11-GA-2012-321693) and the Spanish Ministry of Economy and Competitiveness (Project ref: ECO2012-38134). Gottlieb gratefully acknowledges support from the Becker-Friedman Institute, the Upjohn Institute for Employment Research, and SSHRC. Hémous gratefully acknowledges support from the University Research Priority Program "URPP Equality of Opportunity" of the University of Zurich. Any views expressed are those of the authors and not those of the U.S. Census Bureau. The Census Bureau's Disclosure Review Board and Disclosure Avoidance Officers have reviewed this information product for unauthorized disclosure of confidential information and have approved the disclosure avoidance practices applied to this release. This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 1756 (releases CBDRB-FY22-P1756-R9679 and CBDRB-FY22-P1756-R9845). We are grateful to Frank Limehouse, Shahin Davoudpour, Daniel Polsky, Michael Levin, Leemore Dafny, CIVHC, and the staff at the Chicago Federal Statistical Research Data Center for their generosity in sharing data and facilitating data access. We thank Scott Blatte, Aidan Buehler, Kenta Baron-Furuyama, Oscar Chan, Innessa Colaiacovo, Erin May, Hugh Shiplett, Emilie Vestergaard, Andrew Vogt, and Mingzhe Zhou for outstanding research assistance. Finally, we thank Paul Beaudry, Roland Bénabou, Marianne Bertrand, Austin Goolsbee, David Green, Neale Mahoney, Stefanie Stantcheva, Joachim Voth and numerous seminar participants for valuable comments.

and Heim, 2012). At first glance, this pattern suggests that any explanation for rising inequality—whether globalization, technology, deregulation, or changes to the tax structure—would have to apply to occupations as diverse as bankers, doctors, and CEOs (Kaplan and Rauh, 2013). We argue instead that an increase in income inequality originating within some occupations can spill over into others, driving broader changes in income inequality. Our prime example is physicians, who comprise 13% of the top one percent of wage earners.

Our first contribution is theoretical. We characterize conditions under which an increase in one group's top income inequality increases top income inequality for service providers. This occurs when the services provided are *heterogeneous* in quality and consumers cannot perfectly substitute quality with quantity. These spillovers are geographically local when the services are non-tradable.

We examine our model's predictions in U.S. labor market data. Using a shift-share strategy, we show that a region's top income inequality spills over into top income inequality among physicians, dentists, other medical occupations, and real estate agents. The effect is large, with standardized coefficients of 1.1 to 1.5. We do not observe such spillovers for occupations that do not meet the model's requirements (such as engineers and managers). Using a broader set of occupations and characteristics of occupations, we show that the spillover patterns are consistent with the model's predictions.

Our analysis begins in Section 2 by documenting that the rise in top income inequality is driven primarily by growing income inequality within occupations. We decompose wage income changes from 1980 to 2012 and find that about 70% of the rise in the 99th-to-90th-percentile income ratio is within-occupation.

In Section 3, we develop a theory under which income inequality can spill over from one occupation to another. In our model, widget makers with heterogeneous incomes buy services from doctors who have heterogeneous abilities and thus provide medical services of heterogeneous quality. Consumption of medical services is non-divisible: Each widget maker needs to consume one unit of one doctor's services. In addition, production is not scalable: Each doctor can only serve a fixed number of widget makers. This gives rise to a positive assortative matching mechanism. When both groups' ability distributions have Pareto tails, doctors' income distribution also has a Pareto tail. An (exogenous) increase in income inequality among the widget makers increases relative demand for the services of the highest-ability doctors and increases top income inequality among doctors. Interestingly, increasing inequality of doctors' ability may actually *reduce* doctors' top income inequality. Our results generalize to the case when doctors have a finite positive supply elasticity, and to the case where quantity and quality of medical service are partly (but not perfectly) substitutable. In contrast, with perfect substitution between quality and quantity, any change in widget makers' income distribution only affects the price per unit of quality-adjusted medical service with no consequence for doctors' inequality.

Our baseline model deliberately focuses on local consumption spillovers by considering a single economy and it abstracts from occupational mobility at the top of the income distribution. Our results are robust to allowing for both occupational and geographical mobility of doctors; in both cases, increasing local income inequality of widget makers increases local income inequality for doctors. In contrast, when we allow for trade in medical services, spillovers occur at the national level so that local top income inequality of doctors is independent of local top income inequality.

Section 4 introduces our empirical analysis of inequality spillovers in U.S. local labor markets. We use restricted-access Census and American Community Survey data from 1980 to 2012 to build a panel of labor market areas (LMAs, which are aggregates of commuting zones) and conduct our analysis at this level. Guided by our model, we measure top income inequality as the inverse Pareto parameter for individuals in the top 10% of the local income distribution. We measure inequality for each occupation, *e.g.* physicians, in each region. We then regress local top income inequality among physicians on local top income inequality in the rest of the population.

An OLS regression could suffer from several endogeneity concerns, so we rely on a shift-share strategy. For each LMA, we compute a weighted average of national occupational inequality. The weights correspond to the importance of each occupation in each LMA at the beginning of our sample. In other words, we only exploit the changes in local income inequality that arise from the occupational distribution in 1980 combined with the nationwide trends in occupation-specific inequality. This weighted average serves as our instrument for general inequality in the LMA. We follow the identification assumption of Goldsmith-Pinkham, Sorkin, and Swift (2020), and assume that the occupational shares are uncorrelated with changes in the outcome variable other than through their effect on local top income inequality.

The model predicts local inequality spillovers for occupations providing services that are heterogeneous in quality, non-divisible, and non-tradable. In Section 5, we focus on three high-earning occupations satisfying these criteria: physicians, dentists, and real estate agents. We find positive spillover coefficients, and the effect is economically significant: a 1 standard deviation increase in local top income inequality raises doctors' top income inequality by 1.1 standard deviation, and dentists' and real estate agents' top income inequality by 1.5 standard deviation.

We next estimate spillover coefficients for an additional 27 occupations common in the top 10% of the income distribution. Most of these occupations do not entirely fit the requirements of our theory, and, with a few exceptions, we do not find spillovers for these occupations. In line with our theory, the spillover coefficients are positively correlated with measures of the importance of customer service and of working directly with the public from O*NET and negatively correlated with a measure of tradability. Together, physicians, dentists, and real estate agents already account for 16% of the top 1% of earners. We consider them as poster-child occupations where spillovers are relatively easy to identify, but we hypothesize that spillovers may affect additional occupations even if our empirical strategy fails to detect them: Spillovers may occur at the national level, they may only affect subcategories within an occupation (*e.g.* personal lawyers but not corporate lawyers), or we may simply have too few observations to detect them. Beyond occupations common in the top 10%, we find evidence of spillovers for a range of medical occupations (therapists, veterinarians, etc.).

We conduct numerous robustness checks and extensions on these core results. In one noteworthy exercise, we leverage the financial deregulation of the 70s, 80s, and 90s as one specific policy-induced shock to consumer income inequality, rather than using our shift-share instrument which aggregates many shocks. With this policy-driven instrument, we continue to find inequality spillovers to doctors and dentists.

Finally, we calibrate our model in Section 6 to quantify its implications using New York State data. The calibrated model first shows that shifts in consumer income can explain the observed rise in doctors' income inequality. Second, spillovers to doctors alone amplify a given shock to top inequality by 16%.[1] Third, it quantifies how much spillovers dampen the increase in welfare inequality relative to nominal income inequality for consumers.

This paper contributes to a large literature on the rise in top income inequality and its causes (Piketty and Saez, 2003; Atkinson, Piketty, and Saez, 2011). We build on the "superstars" idea of Rosen (1981), who explains how small differences in talent may lead to large differences

[1]For this calculation, top inequality is measured by the ratio of aggregate income earned by those in the top 1% to aggregate income earned by those in the top 10%.

in income. The key element in his model is the indivisibility of consumption, which leads to a "many-to-one" assignment problem as each consumer only consumes from one performer (singer, comedian, etc.), but each performer can serve a large market (see also Sattinger, 1993).[2] Income inequality among performers increases because technological change or globalization allows the superstars to serve a much larger market—that is, to scale up production.[3] Specifically, let $w(z)$ denote the income of an individual of talent z, $p(z)$ the average price for her services, and $q(z)$ the quantity provided, so $w(z) = p(z)q(z)$. The standard interpretation of "superstars" is that they have very large markets (high $q(z)$). This makes such a framework poorly suited for occupations where output is not easily scalable.

In contrast, we study an assignment model that is "constant-to-one" where superstars are characterized by a high price $p(z)$ for their services. This makes us closer to Gabaix and Landier (2008). They argue that, since executives' talent increases the overall productivity of firms, the best CEOs are assigned to the largest firms. They show empirically that the increase in CEO compensation can be fully attributed to the increase in firms' market size. Grossmann (2007) and Terviö (2008) present models with similar results. Our baseline model with a Cobb-Douglas utility between medical services and the outside good is similar to their production function which is multiplicative in CEO skill and firm productivity, but we focus on a second moment of the income distribution (the Pareto tail instead of the mean) and consider a consumption problem. Importantly, we extend the analysis beyond the baseline model by considering general utility functions, an intensive margin, different entry margins, geographical reallocation, and limited substitutability between quality and quantity.[4]

Our theory offers an amplification mechanism where any shock to top income inequality can spill over to other occupations. The "original" shock may arise from various channels: technological change affecting firm size (Geerolf, 2017), globalization (Bonfiglioli, Crinò, and Gancia, 2018), an increase in innovation (Jones and Kim, 2018; AABB+, 2019), tax system changes (Piketty, Saez, and Stantcheva, 2014), or increased occupational specialization (Edmond and Mongey, 2021). Our theory is agnostic about which explanation matters most for the original rise in income inequality, and focuses on the resulting spillovers to other occupations. Nevertheless, we trace out the spillover effects of a specific shock, namely deregulation in the financial sector as studied by Philippon and Reshef (2012).

Finally, our paper relates to a literature on demand spillovers. In the sociology literature, Wilmers (2017) shows in OLS panel regressions a positive association between wage inequality and dependence on high-income consumers at the industry level. Manning (2004) and Mazzolari and Ragusa (2013) relate the polarization of labor markets to an increase in high-skill workers' demand for low-skill services. Leonardi (2015) argues that high-skill workers also demand relatively more services from other high-skill workers, a pattern that can amplify increases in the skill premium.[5] Importantly, our focus is not the skill premium across occupations, but on changes in top income inequality within an occupation.

The Supplemental Appendix (Gottlieb, Hémous, Hicks, and Olsen, 2026), hereafter SA, contains additional theoretical results, data details, robustness checks, and calibration results.

[2]Adding network effects, Adler (1985) goes further and writes a model where income can drastically differ among artists of equal talents.

[3]Koenig (2023) provides empirical evidence for entertainers using the roll-out of television.

[4]Relatedly, Määttänen and Terviö (2014) build an assignment model for housing, which they calibrate to six U.S. metropolitan areas. They find that the increase in income inequality has led to an increase in house price dispersion (see also Landvoigt, Piazzesi, and Schneider, 2015).

[5]In Buera and Kaboski (2012), structural change leads to a rise in the skill premium as the demand for skill-intensive service increases with income.

2. THE RISE OF WITHIN-OCCUPATION TOP INCOME INEQUALITY

We begin by showing the importance of within-occupation trends in top income inequality. Among workers with positive wage and salary income, the ratio of incomes at the 98th to 90th percentile rose from 1.7 to 2.0 between 1980 and 2012.[6] The ratios also increased for physicians—from 1.5 to 1.7—and for dentists and real estate agents.

To systematically understand the role of occupations, Figure 1 decomposes overall changes in wage income from 1980 to 2012 into within- and between-occupation components. The circle-marked series in Panel (a) shows the overall change in log wage income at each percentile of the distribution. This reproduces the well-known fact that incomes have grown the fastest in the top of the distribution during this time period. We then adapt the within- and between-firm decomposition of SPGB+ (2019) to use occupations instead of firms.[7] This series shows that average log income in the top-percentile bin rose by 0.56 log points during this period, and the square-marked series shows that increases in within-occupation income inequality drove 0.24 log points of this total.

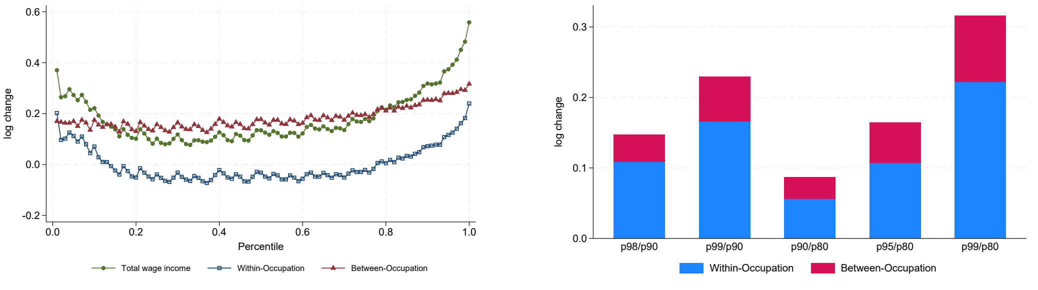

(a) Changes along the income distribution (b) Decomposition for percentile ratios

FIGURE 1.—Changes in wage income from 1980 to 2012: between and within occupations. *Notes:* Panel (a) shows the log change in wage income between 1980 and 2012 for all percentiles of the income distribution (circle-marked). Following the method of SPGB+ (2019), we decompose this change into changes attributable to between-occupation (triangle) and within-occupation (square) changes. Panel (b) shows the contributions of between- and within-occupation effects for top income inequality. p99/p90 is constructed using average log income in the top-percentile and 90th-percentile bins.

Regardless of how we measure top income inequality, inequality within occupations plays a central role. Panel (a) shows that average log income in the top-percentile and 90th-percentile bins rose by 0.56 and 0.32 log points, respectively, implying a 0.24 increase in their difference. Within-occupation factors account for 0.17 ($= 0.24 - 0.07$) of this change, *i.e.* 70% of the total. Panel (b) shows similar patterns for other percentile ratios; 65–75% of the rise in top

[6] We rely on the Decennial Census (1980, 1990, 2000) and the American Community Survey 2010-2014 waves (henceforth 2012). SA Table D.1 shows the corresponding changes for selected occupations. Details on the data are in Section 4 and SA B.1.

[7] To find the effect of within-occupation changes, we hold the average log wage income for occupations fixed at the level of 1980 and only include the changes in the distribution around the averages. For the between-occupation changes we hold the distribution around the average constant, but change the average log wage for occupations. Due to the binning, these two effects don't identically sum to the total change, though in practice the differences are small. See SA B.1 and SPGB+ (2019)'s Online Appendix E for further details.

inequality reflects within-occupation changes. Within-occupation inequality rose most in the top, motivating our focus on the top 10%.[8]

3. THEORY

Motivated by these patterns, we build an assignment model between doctors and their patients to study inequality spillovers across occupations. Section 3.1 presents a special case and Section 3.2 the more general results. Section 3.3 relaxes several assumptions, including allowing for occupational and geographical mobility. Section 3.4 summarizes our empirical predictions.

3.1. *The Cobb-Douglas special case*

We consider an economy populated by two types of agents: widget makers of mass 1 and (potential) doctors of mass $\mu \in (0,1)$.

Production. Widget makers represent the general population. They produce widgets, a homogeneous numeraire good. A widget maker of ability x can produce x widgets. The ability distribution is Pareto with parameter $\alpha_x > 1$ on $x \geq x_{\min}$, such that a widget maker has ability $X > x$ with probability $P(X > x) = \left(\frac{x_{\min}}{x}\right)^{\alpha_x}$. The Pareto parameter, α_x, is an (inverse) measure of the spread of abilities. We treat α_x as exogenous throughout and a fall in α_x captures a general increase in top income inequality.[9] Such a change could arise from globalization or new technology and directly impacts widget makers but not doctors. We set $x_{\min} = \frac{\alpha_x - 1}{\alpha_x}\widehat{x}$ to fix the mean at $\widehat{x}$ when α_x changes.

Doctors produce health services and can each serve λ customers, where we impose $\lambda > \mu^{-1}$ so that there are enough doctors to serve everyone. Potential doctors differ in their ability z, according to a Pareto distribution with shape α_z. They have ability $Z > z$ with probability $P(Z > z) = \left(\frac{z_{\min}}{z}\right)^{\alpha_z}$. All potential doctors can alternatively work as widget makers and produce widgets at some constant ability, which, without loss of generality, we set at $x_{\min}$. (In Section 3.3.2 we instead let an individual's potential ability as a doctor and a widget maker be perfectly correlated.) Unlike Rosen (1981), the ability of a doctor does not change how many patients she can treat. Instead, her skill increases the utility patients get from her care.

Consumption. Widget makers are also doctors' patients. Their preferences over the two goods are represented by the Cobb-Douglas utility function:

$$u(z,c) = z^{\beta} c^{1-\beta}, \tag{1}$$

where c is the consumption of widgets and z is the quality of (one unit of) health care. This quality is equal to the ability of the doctor providing the care. The notion that medical services are not divisible is captured by the assumption that each patient needs to purchase care from exactly one doctor; one cannot substitute quantity for quality. As a result, there need not be a common price per unit of quality-adjusted medical services. More generally, "doctors" here stand in for any occupation which produces non-divisible goods or services for the general

[8] This is consistent with Edmond and Mongey (2021) who, using CPS data, find that residual income inequality has risen for high-skill workers but fallen for low-skill workers. See also Erosa, Fuster, Kambourov, and Rogerson (2025).

[9] Guvenen, Karahan, Ozkan, and Song (2021) show that the top of the income distribution is well-described by a Pareto distribution.

population, including dentists and real estate agents.[10] The Cobb-Douglas utility function eases exposition but can be generalized, as we do in Section 3.2. For simplicity, doctors only consume widgets, so the patients are exclusively widget makers.

3.1.1. *Equilibrium*

Widget makers. Since a widget maker of ability x produces x homogeneous widgets, widget makers' income must be distributed like their ability. The consumption problem of a widget maker of ability x can then be written as:

$$\max_{z,c} u(z,c) = z^{\beta} c^{1-\beta} \text{ subject to } \omega\left(z\right) + c \leq x,$$

where $\omega\left(z\right)$ is the price of medical care from a doctor of ability z. Taking first-order conditions with respect to z and c yields:

$$\omega'\left(z\right) z = \frac{\beta}{1-\beta}\left[x - \omega\left(z\right)\right]. \tag{2}$$

With Cobb-Douglas preferences, no widget maker spends all their income on health care, so equation (2) implies that $\omega(z)$ must be increasing: Higher-ability doctors earn more per patient. Importantly, the non-divisibility of medical services makes doctors "local monopolists" who compete directly only with doctors of slightly higher or lower ability. Therefore $\omega(z)$ need not be linear in z, so there is not generally a fixed price per unit of quality. This makes spillovers from the widget makers' income distribution to doctors' income distribution possible.

As long as the utility function has positive cross-partial derivatives and weakly diminishing marginal utility of consumption, the equilibrium involves positive assortative matching between widget makers' income and doctors' ability (see SA A.1). We denote the matching function as $m\left(z\right)$: a doctor of ability z will be hired by a widget maker whose income is $x = m(z)$.

Doctors. Since there are more doctors than needed, those with lowest z will work as widget makers rather than as physicians. Letting z_c be the ability level of the least able practicing doctor, $m\left(z\right)$ is defined over $[z_c, \infty)$ and $m\left(z_c\right) = x_{\min}$; the worst doctor is hired by a patient with income $x_{\min}$. Market clearing at each z implies:

$$P\left(X > m\left(z\right)\right) = \lambda\mu P\left(Z > z\right), \; \forall z \geq z_c. \tag{3}$$

There are $\mu P\left(Z > z\right)$ doctors with an ability higher than z, each of these doctors can serve λ patients, and there are $P\left(X > m\left(z\right)\right)$ patients whose income is higher than $m\left(z\right)$. With Pareto distributions, equation (3) yields the explicit matching function:

$$m\left(z\right) = x_{\min}\left(\lambda\mu\right)^{-\frac{1}{\alpha_x}}\left(\frac{z}{z_{\min}}\right)^{\frac{\alpha_z}{\alpha_x}}. \tag{4}$$

Intuitively, if top talent is relatively more abundant among doctors than widget makers ($\alpha_z < \alpha_x$), then the matching function is concave. Conversely, it is convex if $\alpha_z > \alpha_x$. We then obtain the cutoff value $z_c = \left(\lambda\mu\right)^{\frac{1}{\alpha_z}} z_{\min}$.

[10]Although we refer to the "quality" of the good, nothing in our model relies on the "high-quality" goods being objectively superior. It is really "quality as perceived by top-earning patients." So a pediatrician who can assuage an anxious parent might have a higher z than one with better diagnostic skills but fewer interpersonal skills. That said, the empirical literature using revealed-preference finds it to be correlated with measures of medical outcomes (*e.g.*, Dingel, Gottlieb, Lozinski, and Mourot, 2026).

Let $w(z)$ denote the income of a doctor of ability z and note that $w(z)=\lambda\omega(z)$ since each doctor provides λ units of health services. As a potential doctor of ability z_c is indifferent between working as a doctor or a widget-maker earning $x_{\min}$, we must have $w(z_c)=x_{\min}$. Combining (2) and (4), we obtain a differential equation that the wage function $w(z)$ must satisfy:

$$w'(z)\,z+\frac{\beta}{1-\beta}w(z)=\frac{\beta}{1-\beta}x_{\min}\left(\frac{\lambda^{\alpha_x-1}}{\mu}\right)^{\frac{1}{\alpha_x}}\left(\frac{z}{z_{\min}}\right)^{\frac{\alpha_z}{\alpha_x}}. \tag{5}$$

Solving equation (5) using the boundary condition at $z=z_c$, we obtain a single solution for the wage profile of doctors. SA A.2.1 demonstrates that this function is:

$$w(z)=x_{\min}\left[\frac{\lambda\beta\alpha_x}{\alpha_z(1-\beta)+\beta\alpha_x}\left(\frac{z}{z_c}\right)^{\frac{\alpha_z}{\alpha_x}}+\frac{\alpha_z(1-\beta)+\beta\alpha_x(1-\lambda)}{\alpha_z(1-\beta)+\beta\alpha_x}\left(\frac{z_c}{z}\right)^{\frac{\beta}{1-\beta}}\right]. \tag{6}$$

As expected, the wage profile $w(z)$ is increasing in doctors' ability z, and $w(z_c)=x_{\min}$. The first term on the right-hand side of (6) dominates for large $\frac{z}{z_c}$ and implies an asymptotic Pareto distribution, so that for large $\frac{z}{z_c}$, we get:

$$w(z)\approx x_{\min}\frac{\lambda\beta\alpha_x}{\alpha_z(1-\beta)+\beta\alpha_x}\left(\frac{z}{z_c}\right)^{\frac{\alpha_z}{\alpha_x}}. \tag{7}$$

Equation (7) shows that the wage schedule at the top is concave in z if $\alpha_z<\alpha_x$; that is, if talent is relatively more abundant among physicians than widget makers. To see why, suppose counterfactually that $\omega(z)\propto z$. Widget makers would then spend a rising share of income on medical services. However, this is in conflict with Cobb-Douglas utility for linear pricing schedules which gives constant spending shares. Hence, this cannot be an equilibrium: demand for high-ability doctors would have to drop, reducing their prices, and resulting in a concave payment schedule. Non-divisibility of medical services is crucial: if high-earning widget makers could simply substitute the services of one doctor of ability z with two doctors of ability $z/2$, a linear pricing schedule would emerge. Conversely, the payment schedule would be convex if talent were relatively scarce among doctors. This wage schedule then guarantees that consumers spend asymptotically a constant share of their income on health care (to see this, combine (4) and (7)). We can thus derive:

PROPOSITION 1: *Doctors' income is asymptotically Pareto distributed with the same shape parameter as for the widget makers. In particular, an increase in top income inequality for widget makers increases top income inequality for doctors. Health expenditures grow proportionately with income: $h(x)\approx\eta x$, where η is a constant.*

To see this, we first define the relevant distribution. For active doctors, the probability that income exceeds w_d is given by $P_{\text{doc}}(W_d>w_d)=\left(\frac{w^{-1}(w_d)}{z_c}\right)^{-\alpha_z}$. Using equation (7), for w_d large enough, we can approximate this distribution as:

$$P_{\text{doc}}(W_d>w_d)\approx\left(\frac{x_{\min}\lambda\beta\alpha_x}{\alpha_z(1-\beta)+\beta\alpha_x}\frac{1}{w_d}\right)^{\alpha_x}. \tag{8}$$

Equation (8) shows that the income of (active) doctors is Pareto distributed at the top. Importantly, in this Cobb-Douglas setting, the shape parameter is inherited from the widget makers, and is independent of α_z, the spread of doctor ability. A decrease in α_x directly translates into a decrease in the Pareto parameter for doctors' income distribution. In other words, the increase in top income inequality spills over from one occupation (the widget makers) to another (doctors). At the top, widget makers' inequality also increases doctors' incomes—a decrease in α_x leads to an increase in $P(W_d > w_d)$ for w_d high enough.[11] Further, a decrease in the mass of potential doctors μ (equivalent to an increase in the mass of widget makers) increases the share of doctors who are active (z_c decreases) and their wages (as $w(z)$ increases if z_c decreases) but has no effect on doctors' top income inequality.

Our results directly generalize to the case where patient income and doctor ability are only asymptotically Pareto distributed and where potential doctors may (but need not) also consume medical services (see details in SA A.3). Our results also hold when doctors' ability distribution has a tail fatter than Pareto. When the tail is thinner than Pareto, a spillover result still holds, although doctors' income is no longer Pareto distributed (see SA A.4).

Taking stock. Proposition 1 establishes the central theoretical result: Changes in widget makers' income inequality translate directly into doctors' income inequality.

3.1.2. *Welfare inequality*

The lack of a uniform quality-adjusted price implies that prices vary along the income distribution. Heterogeneity in consumption patterns implies that people at different points of the income distribution face different price indices (Deaton, 1998). A given increase in income inequality thus translates into a lower increase in welfare inequality. The assignment mechanism implies that as inequality increases, the richest widget makers cannot obtain better health services; in fact, they pay more for health services of the same quality. This mechanism limits the increase in welfare inequality.[12]

To assess this formally, we use a consumption-based measure of welfare. We compute the level of consumption of the homogeneous good $eq(x)$ that, when combined with a fixed level of health quality z_r, gives the same utility to the widget maker as what she actually obtains. That is, we define $eq(x)$ through $u(z_r, eq(x)) = u(z(x), c(x))$. This yields (with a proof in SA A.2.2):

PROPOSITION 2: *For x large enough, welfare $eq(x)$ is Pareto-distributed with shape parameter $\alpha_{eq} \equiv \frac{\alpha_x}{1+\frac{\alpha_x}{\alpha_z}\frac{\beta}{1-\beta}}$. Thus $\frac{d\ln\alpha_{eq}}{d\ln\alpha_x} = \frac{1}{1+\frac{\alpha_x}{\alpha_z}\frac{\beta}{1-\beta}} < 1$, so an increase in widget makers' income inequality translates into a less-than-proportional increase in their welfare inequality. The mitigation is stronger when health services matter more (high β) or when doctors' abilities are more unequal (low α_z).*

3.2. *Generalizing the utility function*

Our results obtain in a much more general case than the Cobb-Douglas utility assumed so far. In SA A.3.3, we show that they generalize to any homothetic utility function that admits positive and finite limits to the elasticity of substitution between health care quality and the homogeneous good when z/c tends to either infinity or 0. For simplicity, we focus here on the

[11]Not all doctors benefit, though, as we combine a decrease in α_x with a decrease in $x_{\min}$ to keep the mean constant. As a result the least able active doctor, whose income is $x_{\min}$, sees a decrease in her income. Had we kept $x_{\min}$ constant so that a decrease in α_x also increases the average widget maker income, then all doctors would have weakly gained.

[12]Moretti (2013) can be viewed as discussing a similar assignment mechanism: high earners locate in high-cost cities, so real inequality is lower than nominal inequality across space.

case when preferences have a constant elasticity of substitution (CES). That is, we replace the utility function of equation (1) with:

$$u(z,c) = \left(\beta z^{\frac{\varepsilon-1}{\varepsilon}} + (1-\beta)\, c^{\frac{\varepsilon-1}{\varepsilon}}\right)^{\frac{\varepsilon}{\varepsilon-1}}, \tag{9}$$

where ε is the elasticity of substitution between health care quality and the homogeneous good. We also only assume that widget makers' and doctors' ability distributions are asymptotically Pareto. With the CES preferences in equation (9), the equilibrium still features positive assortative matching. In SA A.3.2, we show:

PROPOSITION 3: *If either (i) $\varepsilon > 1$ and $\alpha_z^{-1} \leq \alpha_x^{-1}$, or (ii) $\varepsilon < 1$ and $(1-\varepsilon)\,\alpha_z^{-1} < \alpha_x^{-1} \leq \alpha_z^{-1}$, doctors' wages are asymptotically Pareto distributed with shape parameter*

$$\alpha_w^{-1} = \frac{1}{\varepsilon}\alpha_x^{-1} + \left(1 - \frac{1}{\varepsilon}\right)\alpha_z^{-1}, \tag{10}$$

so doctors' top income inequality increases with widget makers' top income inequality. In addition, log health expenditures grow proportionately with log income: $\ln h(x) \approx \left[\left(1 - \frac{\alpha_x}{\alpha_z}\right)\frac{1}{\varepsilon} + \frac{\alpha_x}{\alpha_z}\right]\ln x + \hat{\eta}$, *where $\hat{\eta}$ is a constant.*

Proposition 3 restricts attention to two cases. We argue below that, beyond Cobb-Douglas ($\varepsilon = 1$), these are the two empirically relevant cases. In both cases, doctors' income is asymptotically Pareto distributed with an inverse Pareto parameter that is a linear combination of those for doctors' ability and widget makers' income. An increase in widget makers' top income inequality (α_x^{-1}) leads to an increase in doctors' top income inequality, with a coefficient that has a direct interpretation as the inverse elasticity of substitution between health care quality and the homogeneous good. Health care expenditures increase less than proportionately with income for $\alpha_z \neq \alpha_x$.

To understand the results of Proposition 3 intuitively, and why we restrict attention to certain parameter sets, consider the two cases in turn. First, let medical services and the outside good be substitutes ($\varepsilon > 1$) and let top doctors' skill be relatively scarce ($\alpha_z^{-1} < \alpha_x^{-1}$). In the Cobb-Douglas case ($\varepsilon = 1$), the pricing schedule would be convex (see (7)) and widget makers would spend a constant share of their income on health care. With $\varepsilon > 1$, widget makers reduce their demand for relatively expensive health-care services and health expenditures grow less than proportionately with income, since $\left(1 - \frac{\alpha_x}{\alpha_z}\right)\frac{1}{\varepsilon} + \frac{\alpha_x}{\alpha_z} < 1$. As a result, the wage schedule is less convex than in the Cobb-Douglas case, and the inverse Pareto coefficient for doctors' income distribution, α_w^{-1}, is smaller than α_x^{-1}—but still increasing in α_x^{-1}. In contrast, if $\varepsilon > 1$ and doctors are abundant at the top ($\alpha_x^{-1} < \alpha_z^{-1}$), widget makers would increase their demand for health care services, asymptotically spending nearly all their income on health care, which is counterfactual (and therefore ignored in Proposition 3).

The other case considered in Proposition 3 is when medical services and other goods are complements ($\varepsilon < 1$) and when doctors are relatively abundant in the top: $\alpha_x^{-1} < \alpha_z^{-1}$. With relatively abundant medical services and complementarity between goods, health care expenditures rise less than proportionately with income. Doctors' income is then again asymptotically Pareto distributed with shape parameter $\alpha_w^{-1} < \alpha_x^{-1}$, provided doctors are not too abundant at the top ($(1-\varepsilon)\,\alpha_z^{-1} < \alpha_x^{-1}$).[13]

[13]Equation (10) always holds as long as the right-hand side lies in the interval $(0, \alpha_x^{-1}]$ and covers the Cobb-Douglas case. SA A.3.2 presents the results for the alternative parameter spaces. If $\varepsilon < 1$ and $\alpha_x^{-1} < (1-\varepsilon)\,\alpha_z^{-1}$,

Our empirical analysis suggests that the CES case is more relevant than Cobb-Douglas. First, SA C estimates empirical Engel curves for medical care. We find slopes less than 1, in line with Proposition 3 and in contrast with the Cobb-Douglas case. Second, we will estimate spillover coefficients above 1 (though not always statistically significantly above 1), corresponding to an elasticity $\varepsilon < 1$. Importantly, when $\varepsilon < 1$, a growing spread in doctors' ability (an increase in α_z^{-1}) *reduces* doctors' income inequality (α_w^{-1} decreases). This is because as α_z^{-1} increases, more top doctors compete for patients who are spending a declining share of their income on health care as we move into the tail. In our model, skill-biased technical change for widget makers can be captured by an increase in α_x^{-1}. Similarly, an increase in α_z^{-1} can capture technical change that increases the relative ability of doctors at the top of the distribution. Our results therefore show that this form of technical change for doctors can paradoxically decrease their top income inequality in what will appear to be the most relevant case empirically.

Finally, α_z and α_x are likely to be positively correlated empirically: places with more talent dispersion for widget makers are likely to also have more talent dispersion for doctors. Therefore, without a control for the (unobserved) physician ability distribution, the coefficient of an OLS regression of physician income inequality on widget makers' income inequality would suffer from a downward bias when the true coefficient is above 1—consistent with our findings in Section 5.1.

3.3. *Extensions*

Our model makes several assumptions about preferences, the structure of labor markets and production. To devise appropriate empirical tests for spillovers, we need to establish which assumptions drive the results and which are innocuous. We now show that our results are robust to introducing scalability for medical services, geographical mobility, and occupational mobility for doctors. In contrast, two assumptions are necessary for local inequality spillovers: sufficiently low substitutability between quality and quantity of medical services and non-tradability.

3.3.1. *Scalability*

While doctors' supply was inelastic in the baseline model, we now allow them to increase output at some cost. Specifically, doctors pay effort costs $k\lambda^{1/\varepsilon^S+1}/\left(1/\varepsilon^S+1\right)$ where $k > 0$, $\varepsilon^S > 0$, and λ is the number of patients treated. (Results are identical if doctors pay monetary costs.) Doctors' utility maximization problem immediately implies that their scale depends on the price they can charge as

$$\lambda\left(z\right) = \left(\omega\left(z\right)/k\right)^{\varepsilon^S}, \tag{11}$$

so equation (11) shows that ε^S is the (intensive margin) supply elasticity of medical services. Doctors' income is then given by $w\left(z\right) = \lambda\left(z\right)\omega\left(z\right) = \omega\left(z\right)^{1+\varepsilon^S}/k^{\varepsilon^S}$. The widget makers' consumption problem remains identical and there is still positive assortative matching, though market clearing takes into account that doctors serve different numbers of patients. In SA A.5, we solve the model and prove:

doctors' income is not Pareto distributed and the slope of the Engel curve is asymptotically zero ($\ln h\left(x\right)/\ln x \to 0$) for high incomes. If $\left(\varepsilon - 1\right)\left(\alpha_z^{-1} - \alpha_x^{-1}\right) > 0$, widget makers asymptotically spend all their income on health care and doctors' income is Pareto distributed with shape parameter α_x. SA C shows that the slope of the (log) Engel curve is asymptotically positive and finite, ruling out these two cases.

PROPOSITION 4: *Doctors' income is asymptotically Pareto distributed when: utility is CES with either (i) $\varepsilon = 1$, (ii) $\varepsilon < 1$ and $(1-\varepsilon)\,\alpha_x < \alpha_z < \varepsilon^S + \alpha_x$, or (iii) $\varepsilon > 1$ and $\alpha_z > \varepsilon^S + \alpha_x$. In all cases the shape parameter is*

$$\alpha_w^{-1} = \frac{1+\varepsilon^S}{1+\dfrac{\varepsilon^S}{\varepsilon}\alpha_x^{-1}} \left(\left(1-\frac{1}{\varepsilon}\right)\alpha_z^{-1} + \frac{1}{\varepsilon}\alpha_x^{-1}\right). \tag{12}$$

An increase in widget makers' top income inequality increases doctors' top income inequality ($\frac{d\alpha_w^{-1}}{d\alpha_x^{-1}} > 0$). A higher supply elasticity ε^S increases doctors' top income inequality if and only if $\varepsilon\alpha_x > 1$.

This Proposition shows that the results of Propositions 1 and 3 generalize as long as the supply elasticity of doctors is finite.[14] An increase in widget makers' top income inequality increases doctors' top income inequality, by increasing inequality in both health care prices and quantities supplied (with "weights" $\frac{1}{1+\varepsilon^S}$ and $\frac{\varepsilon^S}{1+\varepsilon^S}$, respectively).

A higher supply elasticity for widget makers would unambiguously increase their income inequality, but a higher ε^S has an ambiguous effect on doctors' top income inequality. Top doctors can disproportionately expand their scale, but the increased supply of high-quality services depresses prices. This price effect is stronger when health care and the homogeneous good are more complementary (ε low), so higher ε^S reduces doctors' income inequality if $\varepsilon < \alpha_x^{-1}$. Therefore, technological change (*e.g.* greater task delegation) that allows top doctors to serve more patients (as in the classic superstar story) need not increase top income inequality. Empirically, a uniform increase in ε^S across regions would not bias our coefficient; but if areas with rising widget makers' inequality also experience larger increases in ε^S, our OLS estimates would be biased upwards when $\varepsilon\alpha_x > 1$.

3.3.2. *Occupational mobility*

We have assumed so far that a potential doctor choosing to work as a widget maker earns the minimum widget maker income, $x_{\min}$. In practice, those succeeding as doctors may have succeeded in other occupations as well (Kirkeboen, Leuven, and Mogstad, 2016). To capture this, we now switch to the opposite extreme and assume perfect correlation between an individual's ability as a doctor and a widget maker. More specifically, we assume that there is a mass of 1 of agents with a unidimensional (Pareto) distribution of skills who decide whether to be doctors or widget makers. While Section 3.3.1 examined the intensive labor supply margin (doctors increasing their scale), this extension studies the extensive labor supply margin (more talented individuals choosing to become doctors). For simplicity, we focus on the Cobb-Douglas case (though similar results can be obtained in the CES case).

SA A.6 shows that Proposition 1 still applies: Doctors' incomes are Pareto distributed with the same coefficient α_x as widget makers, as long as $\lambda^{\alpha_x - 1}\left(\frac{\alpha_z}{\alpha_x}\frac{1-\beta}{\beta}+1\right)^{-\alpha_x} < 1$. This condition ensures that in equilibrium, above a certain threshold, at any given ability some individuals choose to become doctors and others widget makers (otherwise all top individuals choose to be

[14] GPRS+ (2025) estimate a supply elasticity of 0.4 when payments for a doctor's services change. Looking across a broader set of supply responses—beyond those relevant for the Proposition—Clemens and Gottlieb (2014) estimate supply elasticities around 1. There is no contention in the literature that supply is perfectly elastic.

doctors). Intuitively, once individuals have decided their occupation, the ability distribution of doctors is again Pareto, so the results of the baseline model still apply.[15]

3.3.3. *Geographical mobility*

We now extend the baseline model to allow doctors to move across two regions, A and B, of equal size. Medical services are non-tradable and patients are immobile. The regions differ only in the Pareto shape parameter of widget makers' ability distributions, with A more unequal than B ($\alpha_x^A < \alpha_x^B$). Doctors' abilities are Pareto distributed with a common parameter α_z. For simplicity, we focus on the Cobb-Douglas case, though similar results can be derived in the CES case (as well as settings with multiple regions and heterogeneous masses of potential doctors and widget makers).

In autarky, each region mirrors the baseline model: doctors' incomes are asymptotically Pareto distributed with shape parameters matching the local widget maker distribution. With doctor mobility, however, wages $w(z)$ must equalize across regions. Since incomes are higher at the top in region A, high-ability doctors from B migrate to A. As a result, A has a more unequal distribution of doctors' ability than B, but both equilibrium ability distributions are asymptotically Pareto. SA A.7 proves:

PROPOSITION 5: *Once doctors have relocated, the income distribution of doctors in region A is asymptotically Pareto with coefficient α_x^A, and the income distribution of doctors in region B is asymptotically Pareto with coefficient α_x^B.*

Thus, doctors' observed income inequality continues to reflect local consumer inequality even with doctor mobility. Our empirical analysis therefore does not require us to model geographic mobility of doctors explicitly.[16]

3.3.4. *Quality-quantity tradeoff*

Our baseline model assumes that healthcare is non-divisible, *i.e.* quantity cannot substitute for quality of health care. How crucial is this assumption? To answer this, we consider the baseline model with Cobb-Douglas preferences between the outside good and health care consumption but now let health care, H, be a CES aggregate of quality z and quantity q. We assume that a patient must consume health care services of the same quality. We replace the utility function in (1) by

$$u(q,z,c) = H^{\beta} c^{1-\beta} \text{ with } H \equiv \left((1-\gamma)^{\frac{1}{\theta}} q^{\frac{\theta-1}{\theta}} + \gamma^{\frac{1}{\theta}} z^{\frac{\theta-1}{\theta}} \right)^{\frac{\theta}{\theta-1}},$$

where θ measures the elasticity of substitution between quality and quantity. In SA A.8, we show that the equilibrium still features positive assortative matching if $\theta < 1$, and as long

[15]With occupational mobility, note that doctors and widget makers interact through both a demand effect and a labor supply effect. Since the wage level is directly determined by doctors' outside option, one may think that the mechanism which leads to spillovers in income inequality is very different compared to the demand-side mechanism of the baseline model. In SA A.6.1, we split the role of widget makers into two (patients and an "outside option") and show that, with Cobb-Douglas preferences, doctors' income inequality is only driven by their patients' inequality.

[16]Interestingly, the nature of the spillover effect is different with and without mobility: In autarky, local income inequality affects doctors' pay schedule and thereby doctors' local income inequality. With mobility of doctors and many regions, a change in local income inequality does not affect the pay scale for doctors as a function of their ability at the top (except for the most unequal region), but it changes the local *ex post* ability distribution and thereby doctors' local income inequality.

as $\alpha_z > \alpha_x - 1$, doctors' income distribution is asymptotically Pareto distributed with top income inequality spillovers from consumers to doctors (with $\alpha_z < \alpha_x - 1$ doctors' income is bounded). In contrast, if $\theta = 1$, there is no positive assortative matching and doctors' income is proportional to $z^{\frac{\gamma}{1-\gamma}}$.

3.3.5. *Tradability*

Finally, we permit trade in medical services across regions. Since our empirical analysis relies on local spillovers of income inequality, this extension explicitly addresses predictions when services are not sold in a local market. Consider the baseline model of Section 3.1, but with several regions indexed by $s \in \{1, \ldots, S\}$. We allow some patients (a positive share of rich widget makers) to purchase their medical services across regions. The distribution of potential doctors' ability is the same in all regions, and so is the number of patients served per doctor, λ. The other parameters—in particular the Pareto shape parameter of widget makers' income α_x^s—are allowed to differ across regions. It follows that in the top, national income is asymptotically distributed with the lowest α_x^s, that of the most unequal region.[17]

The cost of high-quality health care services must be the same everywhere; otherwise, the widget makers who can travel would go to the region with the cheapest health care. Therefore top talented doctors must all earn the same wage for the same ability. Since national top income inequality for widget makers is $\min_s \alpha_x^s$, doctors' income in all regions is asymptotically Pareto distributed with shape parameter $\min_s \alpha_x^s$ (see details in SA A.9). That is, there are national but not local spillovers.[18]

It is important to recognize that spillovers still exist. An increase in income inequality in the most unequal region continues to determine the income inequality of physicians (in all regions). However, an empirical analysis based on local spillovers will fail to find an effect at least in the tail. Empirically, whether the service provided is "local" (non-tradable) or "non-local" (tradable) will depend on the occupations of interest. We will use these results to guide our empirical analysis.

3.4. *Empirical predictions and institutional details on health care*

To summarize, our model makes the following predictions (where "doctors" represent all occupations that fit the assumptions). (1) High-earning consumers are treated by more expensive doctors. (2) An increase in local inequality will increase local inequality for doctors if they serve the general population directly and their services are non-divisible (or more generally feature limited substitution between quality and quantity). (3) This is true regardless of whether (i) doctors can move across regions and (ii) doctors' ability is positively correlated with the income they would receive in alternative occupations. (4) If medical care is tradable, doctors' top income inequality in each region does not depend on local income inequality, but on national inequality.

While "doctors" are a primary example of an occupation with inequality spillovers, the medical industry in the U.S. is not perfectly described by the flexible price-setting of our model. The government plays a substantial role through Medicare and Medicaid, the insurance sector acts as an intermediary, and there is information asymmetry between patients and doctors. But

[17] If doctors are mobile and medical services are also tradable, the geographic location of agents is undetermined in general, and we would need a full spatial equilibrium model to generate empirical predictions.

[18] Formally, we show in a model with a continuum of agents that the income distribution of doctors is approximately Pareto with shape parameter $\min_s \alpha_x^s$ above a certain cutoff for any positive share of mobile patients. In practice, national spillovers would be negligible if only a small fraction of patients travel.

these institutional intricacies need not inhibit market forces—including our spillovers—from operating. In fact, they may offer a mechanism that implements the forces our model discusses. Providers' negotiations with private insurers generally lead to higher prices in the private market than under public insurance (Clemens and Gottlieb, 2017). Despite asymmetric information, patients often have clear beliefs about who the "best" local doctor in a specific field is, whether or not these beliefs relate to medical skill or health outcomes (Kolstad, 2013, Epstein, 2006, Steinbrook, 2006). Finally, though Dingel, Gottlieb, Lozinski, and Mourot (2026) document patients sometimes traveling for care, the distance sensitivity is high.[19] Therefore, despite these complications, the structure of the health industry may embody enough flexibility to incorporate the local economic pressures implied by our model.

SA C provides evidence in favor of our model's first prediction, namely that there is positive assortative matching between patients' income and medical providers' prices, using medical claims data. We also document that health spending increases with income, using a nationally representative survey. In the rest of the paper, we test our core prediction: within a local geographic market, inequality spills over into occupations that provide non-divisible services with heterogeneous quality.

4. EMPIRICAL STRATEGY TO IDENTIFY SPILLOVERS

This section introduces the main empirical test of our model, using geographical variation in income inequality across LMAs in the U.S. Our goal is to estimate the causal effect of general income inequality on income inequality within a specific occupation.

4.1. *Income data*

We use the Decennial Censuses from 1980, 1990, and 2000 and the 2010-2014 (referred to as "2012") waves of the American Community Survey (ACS) (U.S. Census Bureau, 2000, 2014). We refer to this combined data as Census data. We use 2010-2014 as opposed to 2008-2012 to avoid the immediate aftermath of the Great Recession, which had a large impact on top incomes. We access the restricted-use versions of each which contain a larger sample of respondents and less income censoring than public-use versions.[20] Pre-1980 data come from substantially smaller samples, and we exclude them from the main analysis. SA B.2 discusses the definitions of occupations and geographic locations (Labor Market Areas, LMAs) that we use.

Motivated by our theoretical model, we measure a distribution's top income inequality by its estimated Pareto parameter. Consider a set of observations $\mathcal{N} = \{x_i\}_{i=1}^N$ drawn from a Pareto distribution with two parameters: the minimum value and the Pareto parameter. The maximum likelihood estimate for the minimum value is $\hat{x}_{\min} = \min\{x_i\}_{i=1}^N$ and for the Pareto parameter:

[19] Dingel, Gottlieb, Lozinski, and Mourot (2026) show that in 2017 around four-fifths of medical care is consumed in the same Hospital Referral Region (roughly the size of the Labor Market Areas we use) where it is produced. Our strategy focuses on an earlier period, 1980–2012, where trade in medical services was lower. It also more heavily weights large metropolitan areas, which tend to have more types of care and expert physicians available, so patients from these regions travel for care even less frequently. Regardless, our theoretical analysis suggests that national travel would reduce our estimated regional spillovers.

[20] The detailed Decennial Census long-form surveys 1/6th of the population. Each of the five ACS samples is 3%, so combining them gives a sample of 15%. We adjust incomes for inflation to 2000 dollars. Whereas the publicly available data is censored at around the 99.5th percentile of the overall income distribution, the restricted data has very little censoring. For instance, in New York State only around the top 0.1% is censored. Among physicians the number is well below 0.5%.

$$\hat{\alpha}^{-1} = \frac{1}{N} \sum_{i \in \mathcal{N}} \ln \left(\frac{x_i}{\hat{x}_{\min}} \right). \tag{13}$$

That is, the estimate of the inverse Pareto parameter is the average log distance from the chosen minimum value. So $\hat{\alpha}^{-1}$ is a measure of income inequality even if the distribution is not exactly Pareto. We adjust (13) for the small number of censored observations (see SA B.1). Guvenen, Karahan, Ozkan, and Song (2021) and Jones and Kim (2018) also use α^{-1} as a measure of inequality.

Our focus on top income inequality requires choosing $x_{\min}$. Since the Pareto distribution is a good fit in the top decile, we set $x_{\min}$ as the 90th percentile of the local income distribution of employed adults (age $\geq$ 25) with positive wage income.

Our analysis is demanding of the data: for a given occupation in a given location, computing top income inequality requires many observations. To get enough observations, we use the same $x_{\min}$ to compute all occupation-specific $\hat{\alpha}_o^{-1}$ measures. For example, to compute top income inequality of physicians in New York, we use physicians in the top 10% of New York's overall distribution, which is much more than 10% of New York's physicians because they are high-earners. This approach also implies that we use a consistent slice of the population—those in the upper decile of their area's distribution—when considering different outcome occupations (*e.g.* dentists, real estate agents). This aligns with the motivating observation in Figure 1a that rising within-occupation inequality is concentrated at the top of the distribution.

For our regressions, we will use a balanced panel of the 50 most populous LMAs (as of 1980). These areas tend to have a sufficient number of observations to reliably compute income inequality for the occupations of interest.[21] Including less populated LMAs starts to introduce imbalance in the panel for the smaller occupations of interest. We discuss robustness to these choices below.

4.2. *Income distribution statistics*

We now present summary statistics on income distributions and discuss how well they fit the Pareto distribution. Our main income measure is pre-tax wage and salary income.[22] SA Table D.2 shows descriptive statistics in 2000 for the most common occupations in the top decile of the national income distribution. It reports each occupation's mean income and the share of the top 1%, 5%, and 10% that the occupation represents. Physicians are one of the most common occupations at the top, accounting for 13% of the top 1% (more than the three financial occupations together) and 4% of the top 10%.

We next analyze whether national income distributions are Pareto at the top. With a large number of observations at the national level, we can estimate $\hat{\alpha}^{-1}$ in the top 10% for doctors, dentists, and real estate agents, respectively. We also compute $\hat{\alpha}^{-1}$ for the general population in their top decile. Table I reports the results. $\hat{\alpha}^{-1}$ is mostly larger for the general population than for doctors and dentists which is consistent with the CES model of Section 3.2. To assess the Pareto fit, we compare the ratio of the 98th and 90th income percentiles against the predicted

---

[21] For a Pareto distribution, the maximum likelihood estimator $\hat{\alpha}^{-1}$ has a standard error of $\hat{\sigma} = N^{-1/2}\hat{\alpha}^{-1}$. When $N \geq 20$, the ratio $\hat{\sigma}/\hat{\alpha}^{-1} < 0.23$. In Table VI, we indicate how many LMA-year pairs have an outcome $\hat{\alpha}^{-1}$ calculated using fewer than 20 observations. Using the top 50 most populous LMAs ensures that $N \geq 20$ is nearly always the case for key occupations of interest.

[22] In principle, we would prefer to use all earned income. However, GPRS+ (2025) report that business income is less well measured in Census data than wage income for physicians (their Appendix B.4), so we use wage and salary income for our main measure, and earned income for a robustness check in Section 5.3.

TABLE I

WAGE INCOME: RATIO 98/90: ACTUAL VALUES AND PREDICTED VALUES

	General Population			Physicians			Dentists			Real Estate Agents		
Year	α^{-1}	$\frac{98}{90}$	$\widehat{\frac{98}{90}}$	α^{-1}	$\frac{98}{90}$	$\widehat{\frac{98}{90}}$	α^{-1}	$\frac{98}{90}$	$\widehat{\frac{98}{90}}$	α^{-1}	$\frac{98}{90}$	$\widehat{\frac{98}{90}}$
1980	0.34	1.70	1.72	0.25	1.50	1.50	0.29	1.54	1.61	0.42	1.94	1.97
1990	0.38	1.87	1.85	0.40	1.89	1.90	0.35	1.90	1.76	0.47	2.17	2.12
2000	0.42	2.00	1.96	0.33	1.75	1.71	0.33	1.69	1.70	0.54	2.40	2.37
2012	0.42	1.99	1.96	0.34	1.72	1.72	0.36	1.74	1.78	0.50	2.17	2.22

Note: The inverse Pareto parameter, α^{-1}, is calculated in the top 10% of the relevant national wage income distribution (general population, physicians, dentists, real estate agents). $\frac{98}{90}$ is the ratio of the 98th to 90th percentile of income. $\widehat{\frac{98}{90}}$ is the $\frac{98}{90}$ ratio predicted based on the estimated Pareto parameter as $5^{1/\alpha}$.

98/90 ratio assuming income is Pareto distributed (calculated as $5^{1/\alpha}$). The predicted and actual ratios agree closely, consistent with upper tail incomes being Pareto distributed.

Finally, we assess the extent to which the general and physician income distributions are Pareto at the sub-national level. The Pareto distribution implies a linear relationship between log income and the log of $1-\mathrm{CDF}(\text{income})$. The absolute slope of the relationship is the Pareto coefficient α.[23] We test this prediction in New York State (disclosure rules impede testing it at the LMA level). We bin New York's top decile of earners into 20 equally spaced log income bins. Within each bin, we calculate the share of observations with income in that bin or higher. Figure 2 shows the relationship between the log income bins and log observation shares. Panel 2a shows the relationship for the general population and Panel 2b for physicians specifically. The Pareto fit is excellent for the general population.

For physicians, the distribution is also Pareto in their top 10%, as illustrated by the dashed line. However, the fit worsens lower in the distribution. The α^{-1} estimated on physicians in the top 10% of New York's overall income distribution (as in our regression) is higher than when using the top 10% of physicians. Nevertheless, even when the data are not exactly Pareto, the estimated α^{-1} is a reasonable measure of income inequality.

[23] With N observations of wages drawn from a Pareto distribution, the expected share of observations with a value higher than x, N_x/N, is $(x/x_{\min})^{-\alpha}$. This means $\ln(N_x/N) = -\alpha\ln(x) + \alpha\ln(x_{\min})$.

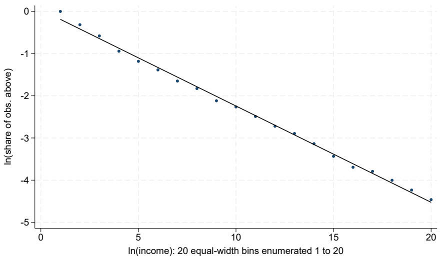

(a) All Occupations

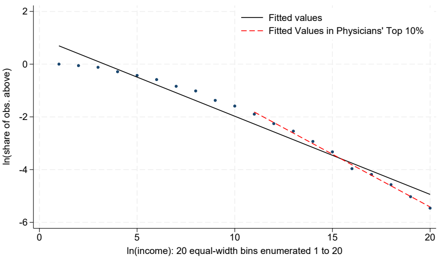

(b) Physicians

FIGURE 2.—Fit of the Pareto Distribution: New York State (2000). *Notes:* This figure shows the quality of fit of the empirical income distribution to the Pareto distribution using observations from New York State in 2000. The left panel shows the top 10% among all occupations, and the right panel shows physicians in the top 10% of the total population. The horizontal axis shows the logarithm of incomes split into 20 equal-width bins. The vertical axis shows the logarithm of $1 - \mathrm{CDF}(\text{income})$. By construction, the line is downward-sloping, and linear if the underlying distribution is Pareto. The lines are linear fits to the binned observations. Panel (b) also shows the linear fit among the subset of bins that correspond to the top 10% of physician earners. Source: Authors' calculations using Decennial data.

SA Figure B.1 shows that, although the estimate of α^{-1} can be sensitive to the choice of sample, the change between years—the source of identifying variation in our regressions—is much less so. Section 5.3 presents robustness checks to ensure that our results are not driven by deviations from the Pareto distribution.

4.3. *Empirical strategy*

Regression framework. We aim to estimate the causal effect of a change in general (population-wide) top income inequality in a region s on the change in top income inequality for a particular occupation o in that region, say, physicians. Let $\alpha^{-1}_{o,t,s}$ be top income inequality for occupation o at time t for geographical area s and $\alpha^{-1}_{-o,t,s}$ be the corresponding value for the general population in s except for o (both computed for individuals in the local top 10%). Let γ_s be a dummy for the geographical area, γ_t a time dummy, and $X_{t,s}$ a vector of controls, including the area's population and average income. The regression for occupation o, at the area-by-year level, is:

$$\alpha^{-1}_{o,t,s} = \gamma_s + \gamma_t + \beta_o \alpha^{-1}_{-o,t,s} + X_{t,s}\delta + \epsilon_{o,t,s}. \tag{14}$$

The coefficient β_o measures the inequality spillover. That is, by how much top income inequality for occupation o increases when top income inequality for the general population increases. We use the Census income data described above to compute $\alpha^{-1}_{o,t,s}$ and $\alpha^{-1}_{-o,t,s}$ for 1980, 1990, 2000 and 2012. We cluster standard errors by LMA and weight LMA-years by the number of persons in occupation o (above $x_{\min}$). With 4 periods and 50 LMAs we have a balanced panel of 200 observations.

Instrument. A natural worry when estimating equation (14) is endogeneity. Even controlling for LMA and year fixed effects, a positive correlation between general income inequality and inequality for a specific occupation might reflect deregulation, tax change, or common local economic trends—rather than a causal effect. Our mechanism itself could generate reverse causality: inequality within the outcome occupation might spill over into other occupations on the right-hand side. And as Section 3.2 explained, unobserved positive correlation between the ability distribution of the occupation of interest and the general population could lead to a downward bias.

To address these concerns, we use a "shift-share" instrument (Bartik, 1991) based on the occupational distribution across geographic areas in 1980. We define:

$$I_{-o,t,s} = \sum_{\kappa \in K_{-o}} \omega_{\kappa,1980,s} \alpha^{-1}_{\kappa,t,-s}, \text{ for } t \in \{1980, 1990, 2000, 2012\}. \tag{15}$$

In equation (15), K_{-o} is the set of most important occupations in the top 10% (excluding o), defined as the union of the 10 most common occupations in the top 10% of each LMA s in 1980. K_{-o} contains around 30 occupations (disclosure rules require rounding to the nearest 10). The weight $\omega_{\kappa,1980,s}$ is the share of occupation κ in LMA s's top 10% in 1980. The shift, $\alpha^{-1}_{\kappa,t,-s}$, is the inverse Pareto coefficient for occupation κ in year t for the entire U.S. excluding LMA s. We estimate equation (14) via 2SLS using $I_{-o,t,s}$ as an instrument for $\alpha^{-1}_{-o,t,s}$. Since K_{-o} does not include all occupations, the weights do not sum to 1. Following Borusyak, Hull, and Jaravel (2022), we control for the sum of the excluded weights interacted with year dummies in all specifications. We explore alternative occupation sets K_{-o} in robustness checks.

The instrument variation is best illustrated by a decomposition of the endogenous variable, income inequality for the general population. Let O_{-o} be the set of all occupations excluding the occupation of interest o. The estimator of the inverse Pareto parameter in (13) implies that we can decompose the common estimate of $\alpha^{-1}_{-o,t,s}$ for a set of occupations O_{-o} as an average of the occupation-specific estimates (where $\alpha_{\kappa,t,s}$ is irrelevant if an occupation does not appear in year t). We exploit this to write our right-hand side measure of inequality as:

$$\alpha^{-1}_{-o,t,s} = \underbrace{\sum_{\kappa \in O_{-o}} \widetilde{\omega}_{\kappa,1980,s} \alpha^{-1}_{\kappa,t,-s}}_{\text{National trends}} + \underbrace{\sum_{\kappa \in O_{-o}} \left(\alpha^{-1}_{\kappa,t,s} - \alpha^{-1}_{\kappa,t,-s}\right) \widetilde{\omega}_{\kappa,1980,s}}_{\text{Local inequality shocks}} + \underbrace{\sum_{\kappa \in O_{-o}} \alpha^{-1}_{\kappa,t,s} \left(\widetilde{\omega}_{\kappa,t,s} - \widetilde{\omega}_{\kappa,1980,s}\right)}_{\text{Changes in occupational composition}}. \tag{16}$$

Here, $\widetilde{\omega}_{\kappa,t,s} \equiv \omega_{\kappa,t,s}/(1-\omega_{o,t,s})$ is occupation κ's share of the upper-tail population excluding occupation o. This normalization ensures that the weights sum to one over O_{-o}. Equation (16) decomposes the change in local inequality into three terms. The first term captures national trends in occupational income inequality, on which we base our instrument (using the part of the sum associated with occupations in $K_{-o} \subset O_{-o}$).[24] Our IV estimation therefore only exploits the changes in local income inequality arising from the 1980 occupational distribution combined with nationwide trends in occupational inequality. Our instrument relies on national trends that reflect shocks exogenous to a given LMA—such as globalization, technological change or deregulation—but which affect LMAs differently depending on occupational composition, leading to a change in α_x in our model. The other two terms in equation (16) capture local LMA changes: the second term reflects local occupational income inequality relative to U.S. trends, and the third captures changes in local occupational distribution. Neither can plausibly be considered exogenous to the occupation of interest o.

We adopt the Goldsmith-Pinkham, Sorkin, and Swift (2020) framework for evaluating shift-share instruments. They show that a sufficient condition for the validity of a shift-share instrument is that the original weights are conditionally exogenous. Our instrument will be valid if the original occupational composition only affects changes in local top inequality for the occupation of interest through changes in local top income inequality (changes, rather than levels,

[24]Mazzolari and Ragusa (2013) use a similar approach to instrument for the *level* of income for high-earners in cities based on pre-sample city-specific occupational distribution and national trends in top income growth.

TABLE II

SUMMARY STATISTICS FOR REGRESSION VARIABLES

	Physicians		Dentists		Real Estate Agents		
	Mean	SD	Mean	SD	Mean	SD	N
α_o^{-1}	0.82	0.12	0.67	0.09	0.50	0.07	200
α_{-o}^{-1}	0.39	0.06	0.40	0.06	0.40	0.06	200
I	0.28	0.02	0.30	0.02	0.29	0.02	200

Note: Summary statistics for the regression variables. α_o^{-1} and α_{-o}^{-1} are the inverse Pareto parameters for the occupation of interest and for the local population excluding the occupation of interest, respectively, in a given LMA×year. Both are based on the top 10% of the wage income distribution in a given LMA×year. I is the instrument: the projected income inequality in a given LMA×year based on the 1980 occupational distribution (see details in text). N is the number of LMA-years.

because we include LMA fixed effects). One concern would be that occupation composition may also affect changes in average incomes, which is why we directly control for this channel. Our identification assumption would be violated if, for instance, in areas which *initially* have more financial managers, physicians are able to get better access to credit *over time*, enabling them to expand their offices, earn higher incomes, and disproportionately so at the top (this corresponds to an increase in ε^S in the model of Section 3.3.1 when $\varepsilon\alpha_x > 1$). We run robustness checks where we exclude financial managers from the instrument.[25] We discuss our shift-share setting further in Section 5.3 and in SA D.2.

Summary statistics for the regression variables. Our first regression results focus on three occupations for which we expect to see spillovers: physicians, dentists, and real estate agents. Table II presents summary statistics for the main regression variables. The average inverse Pareto coefficients differ from the ones reported in Table I particularly for physicians because, as discussed above, we compute them for individuals in the top 10% of the overall local population instead of the top 10% of the occupation.

5. EMPIRICAL SPILLOVER ESTIMATES

This section presents our empirical estimates of spillovers across occupations. In Section 5.1, we first focus on a set of high-earning occupations for which we expect to find spillovers: physicians, dentists and real estate agents. Section 5.2 then considers the 30 most common occupations in the top 10% and demonstrates that spillovers are consistent with the predictions of our model. Section 5.3 carries out several robustness checks. Finally, in Section 5.4 we use financial sector deregulation as a specific exogenous shock to income inequality.

5.1. *Physicians, dentists and real estate agents*

Our central occupations of interest where we expect spillovers are physicians, dentists, and real estate agents. Physicians are a major occupation in the top of the U.S. income distribution (see SA Table D.2). They fit our theory well since they provide a service that is heterogeneous in quality and non-divisible. Dentistry is similar but with fewer intermediaries and regulations. Real estate agents are also common in the top of the U.S. income distribution. Real estate services are non-divisible since home sellers/buyers usually only contract with one real estate

[25]The alternative view on shift-share instruments articulated by Borusyak, Hull, and Jaravel (2022) relies on the exogeneity of the shocks, namely here, trends in occupational inequality. This assumption is likely violated in our case; inequality trends are likely correlated across occupations for reasons other than our spillover through consumption mechanism.

agent.[26] All three occupations also primarily serve their local market, a necessary condition for our empirical test.

Physicians. Table III presents the estimates of equation (14) for physicians. Column (1) shows the OLS regression of physicians' income inequality on general income inequality including year and LMA fixed effects. We find a coefficient of 0.41. This coefficient increases slightly in column (2), where we include controls for LMA population and average wage income among those with positive wage income. Columns (3) and (4) show the first stage regression. The instrument has a reasonable predictive effect on the endogenous variable with an F-statistic around 14. Below all the IV estimates, we show cluster-robust standard errors and then confidence intervals valid when the instrument is weak, using LMMP+ (2025)'s method. Columns (5) and (6) present IV results: Income inequality from the broader population spills over to physician income inequality with an estimated coefficient of 1.54 in the model without controls and 2.29 in the model with controls. Log population has little conditional relationship with physician inequality, whereas log average income predicts lower inequality. SA Figure D.1 visualizes the IV results and demonstrates that they are not driven by outliers.

TABLE III

SPILLOVER ESTIMATES FOR PHYSICIANS

	OLS		First Stage		IV	
Dependent Variable:	α_o^{-1} (1)	α_o^{-1} (2)	α_{-o}^{-1} (3)	α_{-o}^{-1} (4)	α_o^{-1} (5)	α_o^{-1} (6)
α_{-o}^{-1}	0.41	0.58			1.54	2.29
	(0.18)	(0.16)			(0.74)	(0.63)
					[0.28, 3.30]	[1.23, 3.72]
Ln(Average Income)		-0.33		0.06		-0.47
		(0.10)		(0.04)		(0.13)
Ln(Population)		-0.03		-0.03		0.03
		(0.04)		(0.02)		(0.04)
I			4.93	4.44		
			(1.29)	(1.22)		
LMA FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
$(1-\sum_{\kappa\in K_{-o}}\omega_\kappa)\times$ Year FE	Yes	Yes	Yes	Yes	Yes	Yes
N	200	200	200	200	200	200
F-Statistic					14.60	13.21

Note: This table shows OLS and IV regressions of local top income inequality among physicians on top income inequality in the local population. Top income inequality for physicians is measured by the inverse Pareto parameter α_o^{-1} estimated in each LMA among physicians in the top 10% of the local income distribution. Local income inequality is measured by the inverse Pareto parameter α_{-o}^{-1} for the local non-physician population in the top 10% of the local income distribution. I is our instrument and captures the projected income inequality by interacting local occupational composition with national occupation trends in inverse Pareto parameters (see details in text). Column (1) shows the OLS relationship controlling for LMA and year fixed effects and for the share of the LMA's upper tail in 1980 that is not included in the instrument occupations interacted with year indicators. Column (2) adds controls for average income and population among persons with positive income. Columns (3) and (4) show the first stage regressions. Columns (5) and (6) show the IV regressions. N is the number of observations. Observations are weighted by the number of physicians in the top 10% of the local income distribution. The square brackets contain 95% confidence intervals that are robust to weak instruments following LMMP+ (2025).

The magnitude of the spillover relationship is sensible in light of the CES model of Section 3.2. To understand the quantitative implications for the overall increase in physician inequality, we use these estimates and calibrate our model to data from New York State in Section 6.

[26]Furthermore, the fee structure in real estate is often proportional to housing prices (Miceli, Pancak, and Sirmans, 2007) and the increase in the spread of housing prices is consistent with the increase in income inequality (Määttänen and Terviö, 2014).

Alternatively, a model-free approach to interpreting the magnitude of the spillover estimates is to express them in terms of standard deviations (SD) of each variable (as reported in Table II). A 1 SD increase in local income inequality increases local physicians' income inequality by 0.8 SD (from the IV estimate without controls). To give a better sense of these numbers, this would correspond to increasing the income ratio between the 98th and 90th percentiles from 1.73 to 2.00.[27]

Dentists and real estate agents. We show the corresponding results for dentists and real estate agents in Table IV. In both cases, we estimate spillover coefficients greater than 1 in line with the CES model. The IV coefficients with controls are 2.27 for dentists, which is very similar to the corresponding coefficient for physicians, and 1.71 for real estate agents, which is somewhat smaller. In both cases, a one-standard-deviation increase in local income inequality leads to a 1.5-standard-deviation increase in occupational inequality, indicating an economically sizable spillover effect.

TABLE IV

SPILLOVER ESTIMATES FOR DENTISTS AND REAL ESTATE AGENTS

	Dentists				Real Estate Agents			
	OLS		IV		OLS		IV	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
α_{-o}^{-1}	1.08	1.09	2.29	2.27	1.43	1.43	1.69	1.71
	(0.33)	(0.34)	(0.71)	(0.93)	(0.21)	(0.21)	(0.62)	(0.65)
			[1.03, 3.70]	[0.71, 4.47]			[0.65, 2.92]	[0.64, 3.23]
Ln(Avg. Income)		0.08		0.00		0.01		-0.01
		(0.12)		(0.14)		(0.09)		(0.09)
Ln(Pop.)		0.04		0.08		0.02		0.02
		(0.05)		(0.08)		(0.03)		(0.03)
N	200	200	200	200	200	200	200	200
F-Statistic			23.63	13.36			12.72	7.12

Note: This table mimics Table III, except that the outcome occupations are dentists and real estate agents rather than physicians. The square brackets contain 95% confidence intervals that are robust to weak instruments following LMMP+ (2025).

Relationship between IV and OLS. For physicians and dentists, the IV coefficients are much higher than the OLS correlations. What may explain this? First, the CES model of Section 3.2 predicts that an increase in doctors' ability at the top, α_z^{-1}, which could be interpreted as skill-augmenting technical change for doctors, actually reduces top income inequality if the elasticity of substitution $\varepsilon < 1$. That is, technical change that may have increased income inequality for most occupations may actually decrease income inequality for physicians and dentists, countering the spillover effect. This, in turn, implies that a positive unobserved correlation between inequality in local doctors' ability and local consumers' ability would bias the OLS coefficient downward.[28] We discuss this further using our calibrated model in Section 6. Second, there are numerous potential omitted variables. For example, more unequal places might have higher taxes and spend more public money on health care. This would support the incomes of

[27]Formally, physicians' α^{-1} changes by $1.54 \times 0.06 = 0.092$. This is approximately 0.8 times the physicians' SD of 0.12. We then compute the change in the 98th-to-90th-percentile income ratio using the national physicians' $\alpha^{-1} = 0.34$ in 2012 from Table I. For the regression with controls, the effect on physicians is 1.1 SD, which corresponds to increasing the 98th to 90th percentile income ratio to 2.16.

[28]For instance, using (10), the ratio $\text{Cov}\left(\alpha_z^{-1}, \alpha_x^{-1}\right) / \text{Var}\left(\alpha_x^{-1}\right)$ would have to be 1.3 to account for the entire gap between OLS and IV coefficients in columns (2) and (6) in Table III.

worse-off physicians and bias the OLS coefficient downward. Third, due to sampling noise, we estimate inequality in the general population with some error which we expect to bias downward the OLS estimate. For real estate agents, the gap between OLS and IV is much smaller, which in the framework of the CES model, is consistent with a higher elasticity of substitution between the non-divisible service and other goods, ε, and with a smaller spillover coefficient.[29]

5.2. *Other occupations*

We next analyze spillovers for a broader set of occupations. Since Figure 1a shows that rising within-occupation inequality primarily affects the top of the income distribution, we focus on high-earning occupations. First, as a placebo test, we estimate spillover coefficients for three occupations that do not fit our model: financial managers, managers excluding real estate, and engineers. We then examine the 30 most common occupations in the top 10% of the income distribution. We find that (a) spillovers concentrate in occupations fitting our model and (b) spillover coefficients correlate with occupational traits indicating local service provision. Finally, we also find spillovers for a broader class of medical and related occupations that are less common at the top of the income distribution.

TABLE V

SPILLOVER ESTIMATES FOR FINANCIAL MANAGERS, MANAGERS AND ENGINEERS

	Financial Managers				Managers Excl. Real Estate				Engineers			
	OLS		IV		OLS		IV		OLS		IV	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
α_{-o}^{-1}	1.63	1.31	1.12	0.77	0.22	0.28	0.05	-0.02	0.32	0.45	-0.13	-0.09
	(0.47)	(0.35)	(0.97)	(0.92)	(0.09)	(0.09)	(0.12)	(0.15)	(0.11)	(0.13)	(0.26)	(0.30)
			[-0.77, 2.83]	[-1.47, 2.29]			[-0.17, 0.26]	[-0.31, 0.24]			[-0.60, 0.34]	[-0.63, 0.45]
Ln(Avg. Inc.)		0.25		0.29		0.04		0.06		-0.05		-0.02
		(0.14)		(0.16)		(0.03)		(0.03)		(0.04)		(0.05)
Ln(Pop.)		-0.14		-0.16		0.05		0.03		0.07		0.05
		(0.05)		(0.07)		(0.01)		(0.01)		(0.02)		(0.02)
N	200	200	200	200	200	200	200	200	200	200	200	200
F-Statistic			19.92	11.22			22.15	16.49			23.97	27.91

Note: This table shows OLS and IV regressions analogously to Table III but for occupations for which we do not predict spillovers. Columns (1)-(4) look at financial managers, Columns (5)-(8) at managers excluding real estate, and Columns (9)-(12) at engineers. N is the number of observations. The square brackets contain 95% confidence intervals that are robust to weak instruments following LMMP+ (2025).

Placebo occupations. Financial managers, managers (excluding those in real estate), and engineers serve as useful placebos. Workers in these occupations generally do not produce a non-divisible good or service for the local population. Managers and engineers work for firms that produce a variety of goods and services for both the local and national markets.[30] Likewise, financial managers also work for firms: they "plan, direct, or coordinate accounting, investing, banking, insurance, securities, and other financial activities of a branch, office, or department of an establishment" according to the Standard Occupational Classification scheme.

Table V reports the OLS and IV results for these three occupations. The OLS estimates are significant throughout but the IV coefficients are statistically indistinguishable from zero

[29]Proposition 3 also showed that the slope of the Engel curve is less steep when ε is lower. In line with a higher ε for housing than medical services, SA Figure C.1 implies an elasticity of the Engel curve of around 0.29 for medical expenditures, while for real estate, Zabel (2004) finds Engel curve elasticities of 0.64–0.70 for high-income families.

[30]An assignment mechanism may exist for managers but then managers' income inequality would reflect firm size inequality (as in Gabaix and Landier, 2008) rather than *local* income inequality.

and smaller than those for physicians, dentists and real estate agents. The positive OLS estimates and non-significant IV estimates demonstrate that spurious correlation between general inequality and occupational inequality at the local level is likely but that our instrument addresses this concern.

TABLE VI

SPILLOVER ESTIMATES FOR A BROAD SET OF OCCUPATIONS

	OLS (1)	OLS SE (2)	IV (3)	IV SE (4)	F stat (5)	t stat (6)	N Small (7)
Accountants and auditors	0.90	0.12	0.25	0.39	13.26	0.63	0
Airplane pilots and navigators	0.66	0.25	3.54	1.73	6.88	2.04	40
Chief executives, public admin., and legislators	0.77	0.20	0.30	0.60	12.21	0.50	86
Computer programmers	0.43	0.23	0.24	0.43	16.87	0.55	16
Computer systems analysts and computer scientists	0.77	0.23	0.70	0.35	21.87	1.97	7
Dentists	1.09	0.34	2.27	0.93	13.36	2.45	13
Driver/sales workers and truck Drivers	0.84	0.25	1.62	1.12	4.17	1.45	17
Electricians	0.28	0.25	-0.17	1.06	4.44	-0.16	56
Engineers	0.45	0.13	-0.09	0.30	27.91	-0.30	0
Financial managers	1.31	0.35	0.77	0.92	11.22	0.84	0
Financial service sales occupations	1.86	0.30	1.11	1.14	6.07	0.97	21
Insurance sales occupations	1.25	0.17	1.08	0.62	8.07	1.73	2
Lawyers and judges	0.56	0.27	-0.33	0.64	14.21	-0.52	0
Managers of properties and real estate	0.70	0.31	1.32	0.76	15.80	1.75	36
Managers, Excl. Real Estate	0.28	0.09	-0.02	0.15	16.49	-0.10	0
Office supervisors	1.00	0.29	1.41	0.30	15.91	4.70	0
Other financial specialists	0.82	0.28	0.32	0.81	15.49	0.40	0
Personnel, HR, training, and labor rel. specialists	0.67	0.23	0.82	0.60	9.58	1.36	2
Pharmacists	-0.16	0.26	-0.49	0.83	25.01	-0.60	20
Physicians	0.58	0.16	2.29	0.63	13.21	3.63	0
Police and detectives, public service	0.47	0.10	0.30	0.28	27.29	1.08	47
Primary/Secondary School Teachers	0.13	0.28	-0.44	0.64	7.98	-0.68	5
Production supervisors or foremen	0.63	0.18	-0.09	0.75	9.93	-0.12	2
Real estate sales occupations	1.43	0.21	1.71	0.65	7.12	2.64	6
Registered nurses	0.16	0.21	0.71	0.53	27.07	1.36	28
Sales occupations and sales representatives	0.51	0.11	0.28	0.23	8.03	1.20	0
Sales supervisors and proprietors	0.60	0.12	0.53	0.43	20.95	1.23	0
Sales workers	0.69	0.17	0.12	0.43	10.03	0.28	0
Subject instructors, college	0.57	0.11	-0.04	0.28	19.98	-0.15	0
Supervisors of construction work	0.73	0.27	1.27	1.05	6.41	1.21	7

Note: This table shows the OLS and IV coefficients for regressions of local top income inequality for some occupations on top income inequality in the local population excluding that occupation. The occupations shown are the 30 most common occupations in the top 10% of the income distribution. Each row corresponds to the regressions for a given occupation. The variables are defined analogously to Table III and estimates are from the specification with population and average income controls. Columns (1) and (2) show the OLS coefficient and standard error. Columns (3) and (4) show the IV coefficient and standard error. Column (5) is the F statistic for the instrument, column (6) the IV coefficient t statistic. The final column shows the number of LMA-years in which there were fewer than 20 persons used to construct the outcome variable. The total number of LMA-years is 200.

Largest occupations in the top 10% of the income distribution. Table VI reports the results for the 30 most common occupations in the top 10% of the income distribution, including the 6 already considered. Whereas most occupations have a statistically significant positive OLS coefficient, only a small subset have a statistically significant positive IV coefficient. Besides physicians, dentists and real estate agents, we see significant (with $p < 0.1$) positive coefficients for managers of properties and real estate, likely reflecting the same mechanism as for real estate agents, and for insurance sales specialists. The latter often sell their products directly to individuals and, in the occupation characteristics used below, they rank high for customer service, working with the public, and non-tradability. Therefore, these two occupations plausibly fit our spillover mechanism. We also find significant relationships for pilots, computer system analysts, and office supervisors, which may include false positives; with 25 placebo occupations, we would expect 2.5 false positives at $p < 0.1$. Note that for pilots, we have fewer than 20 observations to compute income inequality for 20% of the LMA-years.

The remaining occupations can be classified in three groups. First, those occupations for which our theory does not apply—*e.g.* the placebo occupations already mentioned, computer programmers, and production supervisors. Second, occupations for which our theory may apply but at the national level (such as college professors). Third, occupations (for instance "lawyers and judges") that include both subcategories where our theory should apply (personal attorneys) and where it should not (corporate lawyers). Similarly, financial sales occupations cover both individuals working for firms, for whom our theory may not apply, and personal finance managers for which our theory is more likely to apply, albeit perhaps at the national level.

Relationship with occupation characteristics. Our model predicts a higher local spillover coefficient for occupations where production has to be local, and for those that directly serve the public. To test these predictions, we correlate the IV spillover coefficients with these occupational traits. To quantify local production, we treat Blinder (2009)'s measure of offshorability as an (inverse) measure of the extent to which an occupation serves the local market. It codes 18 out of the 30 occupations as completely not offshoreable. To quantify direct public interaction, we rely on measures from the Occupational Information Network (O*NET; National Center for O*NET Development, 2006) of the level of *customer service* and *working with the public*. SA B.3 gives further details.

Since spillover estimates have varying precision, we divide them by their standard errors; that is, we correlate the t-statistics with occupational traits. Figure 3 shows the results.

Consistent with our model, Panel 3a shows a negative relationship between the spillover and offshorability ($p = 0.07$). Our three focal occupations are non-offshoreable and the non-offshoreable group's average spillover t-statistic is 1.36 compared to only 0.57 in the (at least partly) offshoreable occupations—and this difference is significant at $p < 0.1$. Panels 3b and 3c show positive relationships between the spillover and measures of customer service and working with the public ($p < 0.01$ and $p = 0.11$). SA Table D.3 presents the corresponding regressions, along with the "importance" (rather than "level") of the variables measuring customer service and working with the public, both of which show statistically significant positive relationships.

These patterns support our model: With a few exceptions, we only observe local spillovers for occupations that fit the model's predictions of delivering non-divisible goods or services of heterogeneous quality in the local market. These relationships also buttress our identification strategy; most potential omitted variables would not drive spillover effects for specifically this type of occupation. In particular, neither "keeping up with the Joneses" (Bertrand and Morse, 2016) nor rich individuals sorting into high-amenity locations would imply this spillover coefficient pattern.

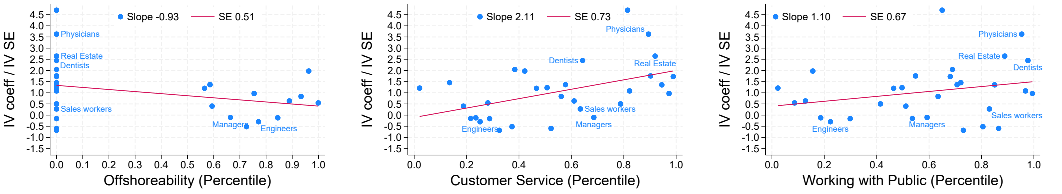

(a) Offshorability (b) Customer service (c) Work w. Public

FIGURE 3.—IV estimates and occupational characteristics. *Notes:* This figure shows the relationship between the t-stat of the spillover coefficients shown in Table VI and three occupation characteristics: a measure of offshorability from Blinder (2009) and two measures from O*NET: the level of "Customer and Personal Service" from Knowledge Requirements and the level of "Performing for or Working Directly with the Public" from Work Activities. The measures are rescaled as percentiles.

Other medical occupations less common in the top. While within-occupation inequality primarily affects top earners, our model does not require occupations receiving spillovers to be common in the top. We test this using four health-related occupation groups less common in the top than physicians or dentists: "veterinarians", "physical and occupational therapists", "psychologists", and "optometrists, podiatrists, and other health-diagnosing occupations". A stacked regression across these groups (allowing for group-specific LMA and year fixed effects) reveals positive spillover effects approximately half the magnitude of our focal occupations (Table VII). Due to sample size limitations, we also present results restricted to the 30 most populous LMAs.[31]

Taking stock. To sum up, we find evidence of spillovers for physicians, dentists, and real estate agents, which together account for a significant share of individuals in the top of the income distribution: 15.8% of the top 1% and 5% of the top 10%. We view this as a lower bound on the applicability of our theory. First, spillovers for occupations with tradable output occur nationally, rather than locally, and therefore would not be captured by our empirical strategy. Second, some occupations common in the top 10%, such as personal lawyers or personal wealth managers, fit the requirements of our model (heterogeneous and non-divisible output) but are not easily identified in Census occupation coding. Moreover, spillovers in theory need not be restricted to occupations in the upper tail, as we show with other medical occupations. With that said, the rise in within-occupation inequality is concentrated at the top of the income distribution, suggesting that spillovers may be particularly prominent (and empirically identifiable) among high earning occupations.

[31]When looking at all 50 LMAs, around 60% of observations have fewer than 20 observations included in our calculation of top income inequality. In the second panel, this falls to 44%.

TABLE VII

SPILLOVER ESTIMATES FOR BROADER MEDICAL OCCUPATIONS

	50 Most Populous LMAs						30 Most Populous LMAs					
	OLS		First Stage		IV		OLS		First Stage		IV	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
α_{-o}^{-1}	0.58	0.54			1.02	1.16	0.05	-0.02			1.18	1.26
	(0.17)	(0.17)			(0.57)	(0.56)	(0.19)	(0.21)			(0.51)	(0.55)
					[-0.01, 2.04]	[0.08, 2.13]					[0.29, 2.25]	[0.33, 2.52]
Ln(Avg. Inc.)		-0.00		0.05		-0.05		0.02		0.04		-0.06
		(0.07)		(0.04)		(0.08)		(0.06)		(0.05)		(0.08)
Ln(Pop.)		-0.05		-0.03		-0.03		-0.05		-0.03		-0.00
		(0.03)		(0.02)		(0.03)		(0.04)		(0.02)		(0.04)
I			4.09	3.97					5.17	4.79		
			(0.80)	(0.99)					(1.17)	(1.32)		
N	800	800	800	800	800	800	500	500	500	500	500	500
F-Statistic					26.15	15.99					19.35	13.03

Note: This table shows regression estimates that "stack" four health-related occupations: Veterinarians, Physical and Occupational Therapists, Psychologists, and the combination of Optometrists, Podiatrists and "Other Health Diagnosing Occupations". Stacking means that for each occupation, we calculate outcome inequality, explanatory inequality, and the instrument. We then stack the data sets together, and run our baseline regression with LMA and Year FEs interacted with indicators for the outcome occupation. N is the number of observations rounded to the nearest integer divisible by 50 as required by disclosure rules. The square brackets contain 95% confidence intervals that are robust to weak instruments following LMMP+ (2025).

5.3. *Robustness checks*

We now discuss robustness checks to alternative hypotheses, specification choices, and inference methods, with corresponding tables in SA D.2.

Earned income. GPRS+ (2025) show that business income is sizable among top physicians. The Census data report wage income, business income, and capital income; we define "earned income" as the sum of the first two. We calculate top income inequality using earned income for physicians, dentists and real estate agents in the same manner as for wage income and replace the dependent variable. We leave the instrument and the independent variables unchanged. The IV coefficients, reported in SA Table D.4, remain similar to those in Table III, though the coefficient for real estate agents declines by around 40%.

Occupational and geographical mobility. Section 3.3 showed theoretically that our spillover mechanism is robust to allowing for either occupational or geographical mobility. We now investigate these issues empirically. First, we restrict attention to physicians that are less occupationally or geographically mobile. Given the cost and time required to enter the medical profession, those older than 35 will have decided to become a medical doctor more than 10 years ago, and cannot easily respond to economic trends over the preceding 10 years. We therefore estimate a regression that restricts attention to doctors older than 35. Similarly for geographical mobility, we consider doctors who have not moved within the past 5 years. Columns (1) and (2) of SA Table D.5 report the results: in both cases, the IV coefficients are a bit higher at 3.1 and 3.3, respectively.

Second, we directly study entry by computing the employment share of the occupation of interest. This measure can vary because of geographical or occupational mobility. In SA Table D.6, we run two sets of regressions, both using our standard IV specification: First, we use employment share as the dependent variable and show that rising inequality does to some extent encourage the entry of new physicians. Nevertheless, when we include it as an additional control in our standard IV regressions with inequality as the outcome we find that our spillover coefficient is virtually unchanged. This is not surprising as entry along the entire ability distribution of physicians would leave physician inequality unchanged while decreasing their mean income. We perform the same analysis with dentists and real estate agents and find that the inclusion of the employment share leaves the spillover coefficient basically unchanged but marginally reduces precision for real estate agents ($p = 0.11$).

Scale. Our theory emphasizes the importance of price inequality in determining doctors' income inequality. Section 3.3.1 demonstrates that our spillover mechanism remains robust even when production scalability is introduced, allowing consumer inequality to raise top performers' relative scale. Yet, technological change that affects the (intensive margin) supply elasticity, ε^S, could affect doctors' income inequality, potentially biasing our IV estimates if correlated with the instrument. Census data lacks direct measures of physician scale. In SA C.3, we use data from the National Ambulatory Medical Care Survey to build a proxy for physicians' scale and find that there is no trend in the standard deviation of that measure between 1991 and 2011. This suggests that, at the aggregate level, there has not been a large shock to ε^S.[32] Similarly, Gilbukh and Goldsmith-Pinkham (2024) find that the share of transactions executed by high commission rate agents has been relatively constant since 2000 (their Appendix Figures F3 and F4).

An increase in dentists' scale is often associated with increased specialization. We proxy for this phenomenon by computing the ratio of dental hygienists to dentists in an LMA. We add this control and our main coefficient is nearly unchanged (SA Table D.5, column (4))—suggesting again that a shock to ε^S is not driving our results.

Medical specialties. Since pay varies across medical specialties and our physician category combines all specialties, compositional effects could influence results. We address this by controlling for specialty shares (detailed in SA D.2). As column (3) of SA Table D.5 shows, the IV coefficient remains virtually unchanged.

Deviations from Pareto distribution. To ensure enough observations and a consistent population slice across occupations, we calculate α_o^{-1} using individuals from the occupation within the top 10% of the general local population. LMA × year × occupation samples are not necessarily Pareto, but this does not undermine identification. First, α_o^{-1} is an inequality measure regardless of distributional form. Our regression coefficient identifies spillovers for that inequality measure provided that our instrument is exogenous—though the IV coefficient may not exactly identify $1/\varepsilon$. Section 6 shows how to calibrate $1/\varepsilon$ using the spillover estimate. Second, as we discussed previously, although α_o^{-1} depends somewhat on the cutoff $x_{\min}$, changes in α_o^{-1}, which we use for identification, are much less sensitive. Third, in SA Table D.7, for our three focal occupations and the three main "placebo" occupations, we control for the share of the occupation in the top 10% of the general population, with little effect on results. Finally, SA Table D.8 shows that calculating α_o^{-1} using only the top 10% of physicians yields a spillover

[32]Further, recall that in our model, an increase in ε^S increases doctors' top income inequality if and only if $\varepsilon\alpha_x > 1$. We can perform a back-of-the-envelope computation to check whether this condition holds. Using (12), we get

$$\beta^{\mathrm{IV}} = \frac{d\alpha_w^{-1}}{d\alpha_x^{-1}} = \frac{\varepsilon^S\left(1 - \alpha_w^{-1}\right) + 1}{\varepsilon^S \alpha_x^{-1} + \varepsilon}.$$

Using the average values in Table I for α_w^{-1} and α_x^{-1}, column (6) of Table III for β^{IV}, and $\varepsilon^S = 0.4$ from GPRS+ (2025), we get $\varepsilon = 0.398$, so $\varepsilon\alpha_x = 1.02$. Taking our estimates and model seriously, parameters are very near the values where top income inequality is independent of ε^S.

coefficient of 1 ($p = 0.12$), or 1.23 ($p = 0.06$) when restricted to the 30 most populous LMAs to reduce measurement noise.[33]

Other sample restrictions. SA Table D.9 reports robustness to other sample choices. First, we use the top 5% of the general population instead of the top 10%. Second, we use 30 or 70 LMAs rather than 50. Third, we construct the instrument using the union of the 13 or 7 (rather than 10) biggest occupations in each LMA. Results are robust to these changes, with one exception: using the 5% cutoff for real estate agents and dentists yields qualitatively similar but less precise coefficients ($p < 0.15$), consistent with having fewer individuals to calculate inequality.

Shift-share robustness checks. Since our identification relies on the exogeneity of the occupational shares, we follow Goldsmith-Pinkham, Sorkin, and Swift (2020) and show the 16 largest (in absolute terms) Rotemberg weights in SA Table D.10. The three occupations with the largest weights are Financial Service Sales (0.35), Financial Managers (0.22), and Airplane Pilots and Navigators (0.21). We exclude, in turn, the five occupations with the highest Rotemberg weight from the IV in our baseline regression for physicians. SA Table D.11 reports the results and shows that the spillover coefficient is very stable.

Alternative inference. Finally, we report alternative inference methods. First, we have shown LMMP+ (2025)-style confidence intervals, which are valid when the instrument is weak. Second, in SA Table D.12 we compute Adão, Kolesár, and Morales (2019) standard errors that account for potential correlation of residual errors across LMAs with similar occupational composition. Neither approach changes the conclusions. In fact, the LMMP+ (2025) confidence intervals are sometimes shorter than under usual inference, which LMMP+ (2025) find is common in many applications.

5.4. *Using financial deregulation as an instrument*

This section investigates spillovers from one specific shock to income inequality: the substantial financial deregulation from the 1970s–1990s.[34] Systematically identifying the source of the original shocks is beyond our scope, but we consider this one in particular since finance occupations are the most influential component of the shift-share instrument according to the Rotemberg weights. Financial deregulation was followed by substantial increases in average income for financial occupations (Philippon and Reshef, 2012). We show that it was also associated with rising income inequality. We use this policy variation, combined with varying concentrations of finance occupations across LMAs, as an exogenous source of variation in local inequality. This exercise requires extending our Census data back to 1960 so we have a pre-treatment period.[35]

[33]In unreported regressions, we ran the same exercise for real estate agents and dentists. We did not find significant spillovers, but note that in both cases we have fewer observations (and much fewer for dentists) than for doctors or for the baseline specification.

[34]Financial deregulation broadly had three components, roughly in this chronological order: removal of restrictions on intra-state branching; removal of restrictions on inter-state banking; erosion of regulatory barriers that separated insurance companies, commercial banking, and investment banks. Many states started intra-state deregulation in the 1970s, most states then experienced inter-state deregulation during the 1980s and early 1990s (Hoffmann and Stewen, 2020). Removal of barriers between commercial banking, investment banking, and insurance occurred throughout the 1990s, culminating in 1999 with the repeal of the Glass-Steagall Act.

[35]Census years 1960 and 1970 contain data quality inconsistencies. For instance, in 1960, some data was lost and later restored, and different county ID schemes were used. In 1970, mail-in forms were introduced for urban areas,

Let $\omega_{f,1970,s}$ denote the share of workers in the top 10 percent of LMA s's wage-and-salary income distribution who work in a finance occupation in 1970, defined as: (i) Financial Managers, (ii) Financial Service and Sales Occupations, and (iii) Other Financial Specialists. The share serves as a measure of exposure to deregulation. We first estimate a continuous difference-in-difference to measure the differential changes in local income inequality (excluding doctors) between areas with high vs. low finance occupation shares relative to 1970:

$$\alpha^{-1}_{-o,t,s} = \theta_s + \theta_t + \sum_{h \neq 1970} \theta_{o,h} \mathbf{1}_{h=t} \omega_{f,1970,s} + v_{o,t,s}, \tag{17}$$

where $\theta_{o,h}$ capture the effect of interest, and θ_s and θ_t are LMA and year fixed effects.

Two identification assumptions are required. First, absent deregulation, areas with more finance would have experienced similar trends in income inequality as areas with less finance. Second, we require that any *direct* effect of finance deregulation on physician inequality is uncorrelated with $\omega_{f,1970,s}$. To illustrate, consider a scenario where deregulation improved physicians' access to credit, which in turn increased physician income inequality. Our identification requires this effect to be uniform across LMAs. If inter-state banking deregulation expanded physicians' credit access more in Columbus than in New York, our spillover estimates would be biased.[36]

Figure 4a visualizes the first stage by plotting estimates of $\theta_{o,h}$. These indicate that consumers' inequality rose more from 1970 to 2000 in areas that had higher finance shares in 1970. A 10 percentage point higher finance share among workers in the local top 10 percent in 1970 predicts a 0.105 increase in α^{-1} for the general population by 1980 and a 0.29 increase by 2000. Panels (b)-(d) plot the reduced form regressions where the outcome in (17) is replaced with income inequality among physicians, dentists, or real estate agents. The finance share predicts changes for physicians and dentists and the reduced form coefficients' pattern approximately tracks the first stage coefficients.

To formalize this into a spillover estimate based on equation (14), we instrument for local inequality with the interaction between $\omega_{f,1970,s}$ and a post-1970 indicator using the data from 1960 to 2012. The results are shown in Table VIII. The spillover estimate for physicians in column (2) is 1.87 and is statistically significant at the 5% level. This is comparable to our baseline estimate of 2.29 from the shift-share design. The estimate for dentists in column (4) is 0.95, which is economically meaningful, but statistically insignificant. On the whole, these results offer some supportive evidence for our main shift-share strategy for doctors and dentists, although less precise. This is not surprising given that this instrument uses substantially less variation than the baseline instrument.

---

and by 1980, nearly all enumeration took the modern form of mailed questionnaires. This motivates our focus on 1980 and later in the main analysis.

[36]Recall that our spillover estimates are largely unchanged when we remove each of the financial occupations (SA Table D.11), so our main results are unlikely to suffer from such a bias. Moreover, such a bias could not give rise to the systematic relationship between spillover coefficients and occupation characteristics shown in Figure 3.

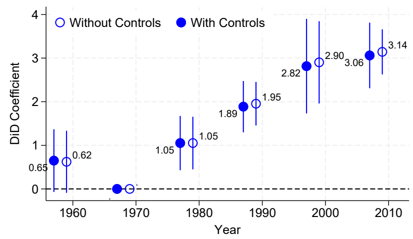

(a) First Stage

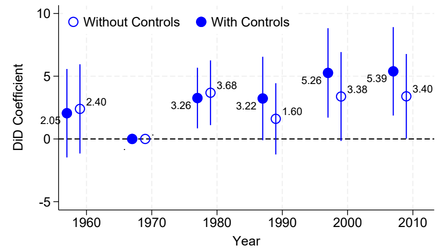

(b) Physicians

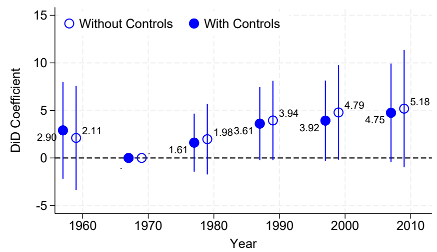

(c) Dentists

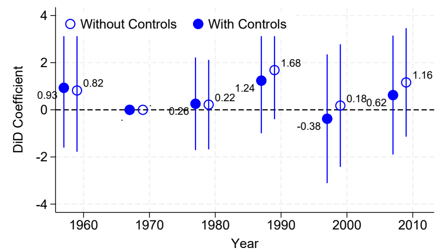

(d) Real Estate Agents

FIGURE 4.—Finance Deregulation Reduced Form. *Notes:* For our baseline set of LMAs, we calculate the share of workers in the local top 10 percent of the wage-and-salary income distribution who work in a finance occupation in 1970: (i) Financial Managers, (ii) Financial Service and Sales Occupations, and (iii) Other Financial Specialists. Denote that share $\omega_{f,1970,s}$. Panel (a) shows estimates of $\theta_{o,h}$ from $\alpha^{-1}_{-o,t,s} = \theta_s + \theta_t + \sum_{h\neq 1970} \theta_{o,h}\omega_{f,1970,s}\mathbf{1}\{h = t\} + v_{o,t,s}$ where o is physicians. The "with controls" version adds ln(population) and ln(average income) controls. Panels (b)-(d) show reduced form estimates from $\alpha^{-1}_{o,t,s} = \pi_s + \pi_t + \sum_{h\neq 1970} \pi_{o,h}\omega_{f,1970,s}\mathbf{1}_{h=t} + e_{o,t,s}$, where the outcome is income inequality among physicians, dentists, real estate agents, respectively, π_t and π_s are year and LMA fixed effects, and $\pi_{o,h}$ are effects of financial deregulation based on pre-existing finance concentrations ($\omega_{f,1970,s}$). 95% confidence intervals are shown.

TABLE VIII

IV SPILLOVER ESTIMATES USING FINANCIAL DEREGULATION INSTRUMENT

	Doctors		Dentists		Real Estate Agents	
	(1)	(2)	(3)	(4)	(5)	(6)
α^{-1}_{-o}	0.93	1.87	1.62	0.95	0.29	0.03
	(0.53)	(0.92)	(1.81)	(1.56)	(0.36)	(0.46)
	[-0.10, 1.95]	[0.75, 3.05]	[-1.87, 5.11]	[-1.65, 3.52]	[-0.37, 0.96]	[-1.23, 1.13]
Ln(Average Income)		-0.39		-0.00		0.08
		(0.16)		(0.13)		(0.09)
Ln(Population)		0.03		-0.16		-0.01
		(0.03)		(0.09)		(0.02)
LMA FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
N	300	300	300	300	300	300
F-Statistic	130.20	26.51	111.00	24.45	45.59	19.91

Note: Spillover estimates using financial deregulation as exogenous variation in local inequality. The second stage is equation (14). The instrument for $\alpha^{-1}_{-o,t,s}$ is the interaction between $\omega_{f,1970,s}$ and an indicator for years after 1970. Data from 1960 to 2012 are used. The square brackets contain 95% confidence intervals that are robust to weak instruments following LMMP+ (2025).

6. QUANTITATIVE INVESTIGATION OF THE SPILLOVER MECHANISM

Our theory establishes that changes in consumers' income inequality translate into changes in doctors' income inequality, particularly in the tail of the distribution. In this section, we solve our model numerically to analyze how these spillovers affect the entire distribution of doctors' incomes and quantify the welfare implications. We use income distributions for both doctors and non-doctors (whom we call "consumers") in New York State for 1980 and 2012.[37]

This exercise serves three purposes. First, we calibrate the model to match our IV coefficient and solve for the endogenous response of doctors' full income distribution to the observed change in consumers' income between 1980 and 2012. We then compare the predicted change to the realized change in doctors' distribution over this period. Second, we quantify the importance of the spillover mechanism for top income inequality in the combined distribution. Third, we quantify how much spillovers dampen the increase in welfare inequality relative to income inequality for the consumers. SA E provides details on the data preparation and calibration.

Data and calibration approach. We observe four empirical income distributions for New York State: doctors and consumers for both 1980 and 2012 (inflation-adjusted to 2000 dollars). Our calibration procedure focuses on matching only the 2012 distributions and the IV coefficient from our regressions. We will then replace *consumers'* 2012 income distribution with their 1980 distribution and compute the resulting physician income distribution implied by the model. This allows us to assess how physician income inequality changed between 1980 and 2012 due solely to spillovers and not other factors such as changes in doctors' ability distribution.

We drop doctors with medical resident-level earnings and the lowest-earning 10% of consumers. We normalize income values by the lowest consumer income in 1980. For each distribution, we fit a flexible kernel to the bottom 90% and a Pareto tail on the top 10%, using the α^{-1} estimated in the top of the data. We use the version of our model from Section 3.2 with CES preferences, in which both consumer income and doctor ability distributions are only restricted to be Pareto in the tail. This more flexible specification allows us to match the full income distribution rather than just its tail behavior. We allow doctors' minimum income $x_{\min}^{\mathrm{dr}}$ to differ from that of consumers to match doctors' significantly higher income.

While we observe consumers' income distribution, doctors' ability distribution, $F_z(z)$, is unobserved and must be calibrated. Our model has four key parameters beyond this ability distribution: (1) the elasticity of substitution between medical services and other goods, ε; (2) the number of patients a doctor can treat, λ; (3) doctors' minimum income, $x_{\min}^{\mathrm{dr}}$; and (4) the preference weight on medical services, β.[38] The elasticity parameter ε is particularly important, since lower values strengthen spillovers. We set $\varepsilon = 0.35$ to target a spillover coefficient of around 1.5 as estimated in the IV panel regression without controls (Table III).[39] We set $\lambda = 175$ to match the observed fraction of doctors in New York State in 2012 (excluding residents), and $x_{\min}^{\mathrm{dr}}$ to match doctors' minimum income. This leaves us with $F_z(z)$ and β to calibrate to match doctors' observed 2012 income distribution. To compute doctors' counterfactual 1980 income distribution, we then keep the calibrated parameters fixed while substituting the 1980 consumer income distribution and $x_{\min}^{\mathrm{dr}}$ value. Table IX summarizes our calibrated parameters.

---

[37] This exercise is most sensible within one geographical market. Disclosure requirements do not permit us to extract data for a specific LMA, so we use New York State instead of New York City.

[38] μ measures the mass of potential doctors and cannot be identified separately from $F_z(z)$, except that, to ensure there are enough doctors to treat everyone, we require $\mu\lambda > 1$.

[39] At first glance, the empirical results in Section 5 along with Proposition 3 would suggest values between $\varepsilon = 1/2.29 = 0.44$ (with controls) and $\varepsilon = 1/1.54 = 0.65$ (without controls). However, since we do not estimate α^{-1} on the extreme tail, the regression-based spillovers somewhat overstate ε.

TABLE IX

CALIBRATED PARAMETERS

$\alpha^{-1}_{\text{cons},1980}$	$\alpha^{-1}_{\text{cons},2012}$	ε	α_z^{-1}	λ	β	$\log(x^{\text{dr}}_{\text{min},1980})$	$\log(x^{\text{dr}}_{\text{min},2012})$
0.34	0.48	0.35	0.51	175	0.01	1.92	2.01

Note: $\alpha^{-1}_{\text{cons}}$ is estimated on the top 10% of the consumers' empirical income distribution. ε is set exogenously to target the IV coefficient (see text). λ is chosen to fit the number of doctors in New York State in 2012. $\log(x^{\text{dr}}_{\text{min}})$ is set to the minimum income of doctors and is scaled by the lowest income of consumers in 1980. β is calibrated along with the ability distribution $F_z(z)$ (including tail α_z^{-1}) to fit the empirical income distribution of doctors in 2012.

Figure 5 shows income distributions for consumers and doctors in the data and in the calibrated model. Panel (a) presents the smoothed consumer income distributions in 1980 and 2012 that we use as inputs in the model. New York State follows the national pattern in Figure 1; income at the 90th percentile rose by 0.40 log points, with smaller increases at lower percentiles. As Table IX shows, α^{-1} rose by 0.14 $(0.48 - 0.34)$ for consumers in New York. Panel (b) shows the corresponding empirical distributions for doctors; income growth was similarly higher at upper percentiles of the doctor distribution; in addition, around 80 percent of doctors (post-medical residency) fall in the top 10 percent of the general income distribution. The 2012 predicted distribution (which the model targets) is visually indistinguishable from the empirical one.

Results. Panel (b) presents our model's counterfactual for the 1980 doctor income distribution, when changing only the consumer income distribution and $x^{\text{dr}}_{\text{min}}$.[40] It also presents the empirical counterpart, which our calibration did not target. Figure 5(b) shows that, in this counterfactual, our model replicates well the overall empirical shift in doctors' income distribution. Panel (c) plots the log income difference between 2012 and 1980 at each percentile of the doctor distribution as predicted by our model, and as observed in the data. At the median, our counterfactual predicts an increase of 0.28 log points compared to the empirical increase of 0.35, while for the 95th percentile, it predicts 0.57 versus the observed 0.51.

While the counterfactual prediction overall matches the empirical change reasonably well, our model over-predicts changes at the very top and under-predicts changes in the middle. This pattern is consistent with doctors' ability becoming more unequal between 1980 and 2012, as discussed in Section 3.2 following Proposition 3. With complementarity between medical services and other goods ($\varepsilon < 1$), a rise in doctors' ability inequality (higher α_z^{-1}) dampens their income inequality. Our model can therefore match the data if we also let doctors' ability distribution change between 1980 and 2012. SA Figure E.2 shows the 1980 and 2012 doctors' ability distributions necessary to perfectly match both years' doctor income distributions. The 2012 ability distribution has a fatter tail than the calibrated 1980 ability distribution. Specifically, α_z^{-1} needs to rise from 0.35 to 0.51—a change comparable to the 0.14 increase in consumers' α^{-1}—for our model to match the data perfectly in both years.

[40]Setting $\varepsilon = 0.35$ ensures that the spillover coefficient is in line with our IV estimate of 1.5. In this counterfactual analysis, the spillover coefficient, computed as the ratio of (1) the change in α^{-1} for doctors between the (empirical) 2012 distribution and the counterfactual 1980 distribution and (2) the empirical change in α^{-1} for consumers (see details in SA E), is in line with our IV estimate. In this exercise, whether we fix $x^{\text{dr}}_{\text{min}}$ or change it to its 1980 value only impacts the very bottom of the predicted distribution.

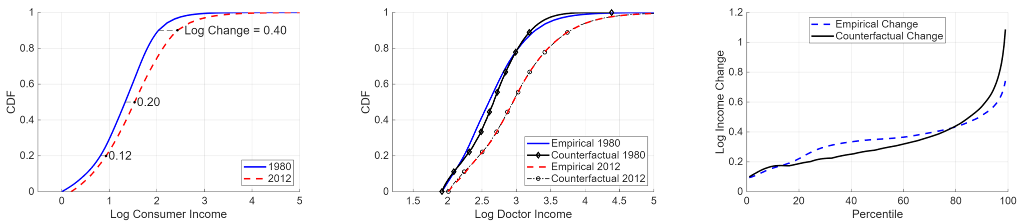

(a) Consumer income distribution (b) Doctor income distribution (c) Change in Doctors' Income

FIGURE 5.—CDFs of income for consumers and physicians in New York State. *Notes:* Panel (a) is the consumers' income distribution. Panel (b) shows both the model and empirical distributions of doctors' income in 1980 and 2012. Parameters are calibrated to match 2012 exactly. "Counterfactual 1980" uses these parameters to compute a counterfactual distribution in 1980 by only changing the consumers' income distribution and $x^{\mathrm{dr}}_{\mathrm{min}}$, when ε is set to target the spillover coefficient from the IV estimate. Panel (c) shows empirical and counterfactual changes in log income for doctors at each percentile of the doctors' income distribution. All values are normalized by the lowest consumer income in 1980.

Spillover effects and within-top income inequality. Second, the calibration enables us to assess the quantitative importance of doctors' income spillovers for top income inequality overall (i.e., for the overall income distribution combining doctors plus consumers). To do this, we compute top income inequality between 1980 and 2012 under two scenarios for doctors' income distribution in 1980. The first scenario is the counterfactual exercise described above where spillovers influence doctors' income distribution. The second scenario is an alternative counterfactual for 1980 that eliminates all changes in doctors' income inequality, including changes driven by the spillover mechanism.[41] We then combine each of the resulting doctors' income distributions with the empirical consumer income distribution to get the combined income distribution in 1980 for both scenarios.

Table X reports the empirical share of income within the top 10% going to various subgroups of the top earners, along with the top 10%'s share of all income. The first entry in column 1 shows that the top 10% earned 34.4% of total income in 2012, with the top 1% capturing 29.9% of income within the top 10% (column 1, row 3). Columns 2 and 3 then show comparable statistics under our two counterfactual scenarios. The spillover model predicts an 8.1 percentage point increase in the share of the top 10% that accrues to the top 1% between 1980 and 2012 (column 2, row 3), compared to only 7.0 points in the alternative without spillovers (column 3, row 3). This substantial difference exists despite doctors comprising only 4.6% of workers in the top 10% of the overall distribution.[42]

To illustrate the potential broader impact of spillovers, columns 4 and 5 scale up the population experiencing spillovers. Together real estate occupations (agents and managers) and dentists are around half as numerous as physicians, and other occupations experience spillovers as well. For illustrative purposes, we double the share of "physicians" in the economy, representing a larger set of occupations subject to spillovers (henceforth, spillover occupations).[43]

[41]In this second scenario, doctors' incomes increase by a constant proportion, set equal to the average proportional increase of the first scenario.

[42]To focus on the role of spillovers, we use our baseline model here where the ability distribution for doctors is held constant at its 2012 value, so that the model does not reproduce the empirical measures of income inequality for the joint distribution in 1980. For context, the share of the top 10% that accrues to the top 1% rose by 7.6 percentage points in the data.

[43]Formally, we double the distribution of potential doctors and halve λ to capture the larger mass of spillover occupations. We recalibrate the model, which leaves $F_z(z)$ unchanged and doubles β.

The model without spillovers now has a larger group with no within-group increase in income inequality and predicts a 6.0 percentage point increase in the top 1% share (column 5, row 3). The model with spillovers predicts an 8.0 percentage point increase (column 4, row 3)—a difference of more than 30%. So even when spillover occupations are a relatively small share of the top income distribution, they can generate substantial additional within-top inequality. To streamline exposition, for the remainder of this section we only focus on the case with this broader category of "spillover occupations".

TABLE X

SPILLOVER EFFECTS AND WITHIN-TOP INCOME INEQUALITY

		Spillovers only to doctors (4.6% of top 10%)		Spillovers to doctors + others (9.2% of top 10%)	
	2012	Δ with Spillover	Δ without Spillover	Δ with Spillover	Δ without Spillover
Inequality Measure	(%)	(percentage points)		(percentage points)	
	(1)	(2)	(3)	(4)	(5)
Top 10%'s Share of All Income	34.4	8.1	7.9	8.1	7.7
Top 5%'s Share of Top 10%'s Income	69.7	6.1	5.6	5.8	4.9
Top 1%'s Share of Top 10%'s Income	29.9	8.1	7.0	8.0	6.0
Top 0.1%'s Share of Top 10%'s Income	8.9	4.3	3.6	4.4	3.1

Note: Income inequality measures are for the whole distribution (consumers + doctors). Column (1) shows empirical 2012 measures for New York State. Columns (2) and (4) show model-predicted changes from 1980 to 2012 with spillovers; columns (3) and (5) show the analogous changes with spillovers shut off by keeping the shape of doctors' income distribution constant at its 2012 level but shifting the distribution down by the same log mean. Columns (2)–(3) treat doctors as the only spillover occupation (4.6% of the top 10%), whereas columns (4)–(5) double the size of spillover occupations.

Nominal and real income inequality. Our model showed that, with spillovers, increases in nominal income inequality overstate changes in consumers' welfare inequality. We now quantify that point. For each percentile of the consumer distribution, we calculate the equivalent variation—the extra income needed in 1980 to reach the same utility level as in 2012. We first calculate consumers' utility realized at each percentile, p, of consumer income, $v_p^t(x_p^t) = u(z^t(x_p^t), x_p^t - \omega^t(z^t(x_p^t)))$ for $t = 1980, 2012$, where x_p^t is the income of a consumer at percentile p in year t and $\omega^t(z^t(x_p^t))$ is the spending on medical services at the equilibrium choice of doctor quality $z^t(x_p^t)$ in year t. We use again our baseline model where doctors' ability distribution is fixed at its 2012 value. We then calculate the equivalent variation, EV_p, as the extra income required in 1980 to bring an agent at percentile p to the utility in 2012; that is,

$$v_p^{1980}(x_p^{1980} + EV_p) = v_p^{2012}(x_p^{2012}). \tag{18}$$

Figure 6(a) shows both the observed change in real income and the EV measure defined by equation (18). Panel (b) shows the difference between these two. For low-income consumers the difference is below 0.01 log points. But for those above the 80th percentile, the log difference between income and welfare changes grows to $0.02 - 0.04$. This difference reflects the disproportionate increase in prices for high-quality services faced by high-income consumers. Although the share of spending on services provided by spillover occupations declines with income, this effect is quantitatively minor compared to the price effect.

The spillover effect implies that conventional measures of income inequality using a common price deflator overstate the increase in welfare inequality. While the measured 90/50 income ratio rose by 0.20 log points, the corresponding welfare ratio increased by 0.18 log points.

These two exercises show that, although spillovers increase top income inequality, they also imply that changes in income inequality overstate changes in welfare inequality. SA E brings these two elements together and shows that, compared to the scenario where doctors all experience the same proportional income change, spillovers reduce welfare inequality outside of the top 1% of the combined income distribution, but increase it in the very top.

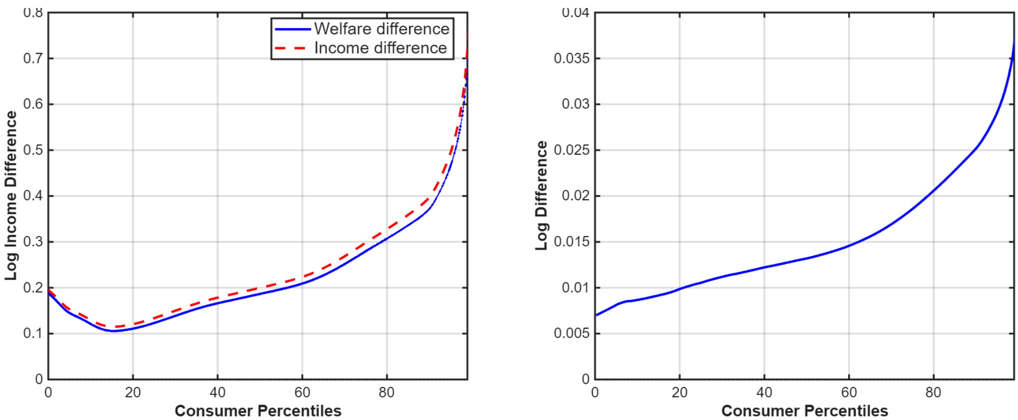

(a) Welfare and Income Changes

(b) Difference between welfare and income changes

FIGURE 6.—Changes in income and welfare for consumers, 1980 to 2012. *Notes:* In Panel (a), "Income difference" is the log change in real income (with a common CPI) for each percentile of consumers and "Welfare difference" is the extra income required in 1980 to reach the utility for the same percentile in 2012. Panel (b) shows the gap between Income difference and Welfare difference shown in Panel (a).

7. CONCLUSION

This paper documents that the majority of the increase in top income inequality in the U.S. is within occupations. We develop a new theoretical framework where an increase in top income inequality in one occupation can spill over through consumption to other occupations that provide non-divisible services directly to customers, such as physicians, dentists and real estate agents. We show empirically that changes in local income inequality do indeed spill over to these occupations, with standardized coefficients ranging from 1 to 1.5. Calibrating our model to New York State, we find that spillovers can account for the entire increase in income inequality for physicians.

Our analysis suggests that the increase in top income inequality observed across most occupations since 1980 may not require a common explanation. Increases in inequality for bankers or CEOs due to deregulation or globalization may have spilled over to other high-earning occupations, increasing top income inequality broadly. While we have emphasized positive results, the theory has an important normative implication: Increasing inequality in the prices of non-divisible services implies that welfare inequality does not rise as much as nominal income inequality. Similar spillover effects could exist for scarce goods such as luxury wine, or sports teams, a question we leave to future research.

REFERENCES

Adão, Rodrigo, Michal Kolesár, and Eduardo Morales (2019), "Shift-share designs: Theory and inference." *The Quarterly Journal of Economics*, 134 (4), 1949–2010. [29]

Adler, Moshe (1985), "Stardom and talent." *American Economic Review*, 75 (1), 208–212. [4]

[AABB+] Aghion, Philippe, Ufuk Akcigit, Antonin Bergeaud, Richard Blundell, and David Hémous (2019), "Innovation and top income inequality." *Review of Economic Studies*, 86 (1), 1–45. [4]

Atkinson, Anthony B., Thomas Piketty, and Emmanuel Saez (2011), "Top incomes in the long run of history." *Journal of Economic Literature*, 49 (1), 3–71. [3]

Bakija, Jon, Adam Cole, and Bradley Heim (2012), "Jobs and income growth of top earners and the causes of changing income inequality: Evidence from u.s. tax return data." Working paper. [1]

Bartik, Timothy J. (1991), *Who Benefits from State and Local Economic Development Policies?*, W.E. Upjohn Institute for Employment Research, Kalamazoo, MI. [19]

Bertrand, Marianne and Adair Morse (2016), "Trickle-down consumption." *The Review of Economics and Statistics*, 98 (5), 863–879. [25]

Blinder, Alan S. (2009), "How many US jobs might be offshorable?" *World Economics*, 10 (2), 41–78. [25]

Bonfiglioli, Alessandra, Rosario Crinò, and Gino Gancia (2018), "Betting on exports: Trade and endogenous heterogeneity." *The Economic Journal*, 128 (609), 612–651. [4]

Borusyak, Kirill, Peter Hull, and Xavier Jaravel (2022), "Quasi-experimental shift-share research designs." *Review of Economic Studies*, 89 (1), 181–213. [19, 20]

Buera, Francisco and Joseph Kaboski (2012), "The Rise of the Service Economy." *American Economic Review*, 102 (6), 2540–2569. [4]

Clemens, Jeffrey and Joshua D. Gottlieb (2014), "Do physicians' financial incentives affect treatment patterns and patient health?" *American Economic Review*, 104 (4), 1320–1349. [12]

Clemens, Jeffrey and Joshua D. Gottlieb (2017), "In the shadow of a giant: Medicare's influence on private physician payments." *Journal of Political Economy*, 125 (1), 1–39. [15]

Deaton, Angus (1998), "Getting prices right: What should be done?" *Journal of Economic Perspectives*, 12 (1), 37–46. [9]

Dingel, Jonathan I., Joshua D. Gottlieb, Maya Lozinski, and Pauline Mourot (2026), "Market size and trade in medical services." Working Paper 31030, National Bureau of Economic Research. [7, 15]

Edmond, Chris and Simon Mongey (2021), "Unbundling labor." Mimeo. [4, 6]

Epstein, Andrew J (2006), "Do cardiac surgery report cards reduce mortality? Assessing the evidence." *Medical Care Research and Review*, 63 (4), 403–426. [15]

Erosa, Andrés, Luisa Fuster, Gueorgui Kambourov, and Richard Rogerson (2025), "Labor market polarization and inequality: A Roy model perspective." Working Paper 33687, National Bureau of Economic Research. [6]

Gabaix, Xavier and Augustin Landier (2008), "Why has CEO pay increased so much?" *The Quarterly Journal of Economics*, 123 (1), 49–100. [4, 23]

Geerolf, François (2017), "A theory of Pareto distributions." Working paper. [4]

Gilbukh, Sonia and Paul Goldsmith-Pinkham (2024), "Heterogeneous real estate agents and the housing cycle." *The Review of Financial Studies*, 37 (11), 3431–3489. [28]

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift (2020), "Bartik instruments: What, when, why, and how." *American Economic Review*, 110 (8), 2586–2624. [3, 19, 29]

[GPRS+] Gottlieb, Joshua, Maria Polyakova, Kevin Rinz, Hugh Shiplett, and Victoria Udalova (2025), "The earnings and labor supply of U.S. physicians." *Quarterly Journal of Economics*, 140 (2), 1243–1298. [12, 16, 27, 28]

Gottlieb, Joshua D., David Hémous, Jeffrey Hicks, and Morten Olsen (2026), "Supplement to 'The Spillover Effects of Top Income Inequality'." Econometrica Supplemental Material. [4]

Grossmann, Volker (2007), "Firm size, productivity, and manager wages: A job assignment approach." *The B.E. Journal of Theoretical Economics*, 7 (1), Article 8. [4]

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song (2021), "What do data on millions of U.S. workers reveal about lifecycle earnings dynamics?" *Econometrica*, 89 (5), 2303–2339. [6, 16]

Hoffmann, Mathias and Iryna Stewen (2020), "Holes in the dike: The global savings glut, U.S. house prices, and the long shadow of banking deregulation." *Journal of the European Economic Association*, 18 (4), 2013–2055. [29]

Jones, Charles and Jihee Kim (2018), "A Schumpeterian model of top income inequality." *Journal of Political Economy*, 126 (5), 1785–1826. [4, 16]

Kaplan, Steven and Joshua Rauh (2013), "It's the market: The broad-based rise in the return to top talent." *Journal of Economic Perspectives*, 27 (3), 35–56. [2]

Kirkeboen, Lars J, Edwin Leuven, and Magne Mogstad (2016), "Field of study, earnings, and self-selection." *The Quarterly Journal of Economics*, 131 (3), 1057–1111. [12]

Koenig, Felix (2023), "Technical change and superstar effects: Evidence from the rollout of television." *American Economic Review: Insights*, 5 (2), 207–223. [4]

Kolstad, Jonathan T (2013), "Information and quality when motivation is intrinsic: Evidence from surgeon report cards." *American Economic Review*, 103 (7), 2875–2910. [15]

Landvoigt, Tim, Monika Piazzesi, and Martin Schneider (2015), "The housing market(s) of San Diego." *American Economic Review*, 105 (4), 1371–1407. [4]

[LMMP+] Lee, David S., Justin McCrary, Marcelo J. Moreira, Jack R. Porter, and Luther Yap (2025), "What to do when you can't use '1.96' confidence intervals for IV." Working Paper 31893, National Bureau of Economic Research. [21, 22, 23, 27, 29, 31]

Leonardi, Marco (2015), "The effect of product demand on inequality: Evidence from the United States and the United Kingdom." *American Economic Journal: Applied Economics*, 7 (3), 221–247. [4]

Määttänen, Niku and Marko Terviö (2014), "Income distribution and housing prices: An assignment model approach." *Journal of Economic Theory*, 151, 381–410. [4, 21]

Manning, Alan (2004), "We can work it out: The impact of technological change on the demand for low-skill workers." *Scottish Journal of Political Economy*, 51 (5), 581–608. [4]

Mazzolari, Francesca and Giuseppe Ragusa (2013), "Spillovers from high-skill consumption to low-skill labor markets." *Review of Economics and Statistics*, 95 (1), 74–86. [4, 19]

Miceli, Thomas J., Katherine A. Pancak, and Clemon F. Sirmans (2007), "Is the compensation model for real estate brokers obsolete?" *Journal of Real Estate Finance and Economics*, 35 (1), 7–22. [21]

Moretti, Enrico (2013), "Real wage inequality." *American Economic Journal: Applied Economics*, 5 (1), 65–103. [9]

National Center for O*NET Development (2006), *O*NET 10.0 Database*, U.S. Department of Labor, Employment and Training Administration (USDOL/ETA). Retrieved from https://www.onetcenter.org/db_releases.html. [25]

Philippon, Thomas and Ariell Reshef (2012), "Wages and human capital in the U.S. finance industry: 1909–2006." *Quarterly Journal of Economics*, 127 (4), 1551–1609. [4, 29]

Piketty, Thomas and Emmanuel Saez (2003), "Income inequality in the United States, 1913–1998." *Quarterly Journal of Economics*, 118 (1), 1–39. [3]

Piketty, Thomas, Emmanuel Saez, and Stefanie Stantcheva (2014), "Optimal taxation of top labor incomes: A tale of three elasticities." *American Economic Journal: Economic Policy*, 6 (1), 230–271. [4]

Rosen, Sherwin (1981), "The economics of superstars." *American Economic Review*, 71 (5), 845–858. [3, 6]

Sattinger, Michael (1993), "Assignment models of the distribution of earnings." *Journal of Economic Literature*, 31 (2), 831–880. [4]

[SPGB+] Song, Jae, David Price, Fatih Guvenen, Nicholas Bloom, and Till von Wachter (2019), "Firming up inequality." *The Quarterly Journal of Economics*, 134 (1), 1–50. [5]

Steinbrook, Robert (2006), "Public report cards—cardiac surgery and beyond." *New England Journal of Medicine*, 355 (18), 1847–1849. [15]

Terviö, Marko (2008), "The difference that CEOs make: An assignment model approach." *American Economic Review*, 98 (3), 642–668. [4]

U.S. Census Bureau (2000), "Decennial Census (SAP 690) [Restricted-Use Data]." URL https://www.researchdatagov.org/product/690. Restricted-use data accessed via Federal Statistical Research Data Centers: 1960, 1970, 1980, 1990, and 2000. [15]

U.S. Census Bureau (2014), "American Community Survey (SAP 691) [Restricted-Use Data]." URL https://www.researchdatagov.org/product/691. Restricted-use data accessed via Federal Statistical Research Data Centers: Vintages: 2010–2014 (5-year ACS sample). [15]

Wilmers, Nathan (2017), "Does consumer demand reproduce inequality? High-Income consumers, vertical differentiation, and the wage structure." *American Journal of Sociology*, 123 (1), 178–231. [4]

Zabel, Jeffrey (2004), "The demand for housing services." *Journal of Housing Economics*, 13 (1), 16–35. [23]