

NBER WORKING PAPER SERIES

RE-EXAMINING GEOGRAPHIC VARIATION IN HEALTH AND HEALTH CARE

Amy Finkelstein
Matthew Gentzkow

Working Paper 34682
<http://www.nber.org/papers/w34682>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2026

We are grateful to the National Institute of Aging (Finkelstein, R01-AG302449) and the Peterson Center on Healthcare for financial support. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w34682>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2026 by Amy Finkelstein and Matthew Gentzkow. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Re-Examining Geographic Variation in Health and Health Care
Amy Finkelstein and Matthew Gentzkow
NBER Working Paper No. 34682
January 2026
JEL No. H51, I1, I11

ABSTRACT

A large literature has documented widespread variation in health care spending per capita across areas of the United States without correspondingly better health outcomes. Recent work has used mover designs to estimate the causal impact of place on both health care spending and mortality. In this paper, we investigate whether places that increase health care spending also tend to be places that increase health. We find that they do not and discuss the implications.

Amy Finkelstein
Massachusetts Institute of Technology
Department of Economics
and NBER
afink@mit.edu

Matthew Gentzkow
Stanford University
Department of Economics
and NBER
gentzkow@stanford.edu

Re-Examining Geographic Variation in Health and Health Care

By AMY FINKELSTEIN AND MATTHEW GENTZKOW*

Mention “Miami and Minneapolis” or “McAllen and El Paso, Texas” to any health economist active over the last quarter century and you are likely to be greeted with both recognition and disagreement. These city pairs are shorthand for a puzzle, documented extensively by researchers at Dartmouth (e.g., Fisher et al., 2003a,b) and popularized in the *New Yorker* by Atul Gawande (Gawande, 2009), that has intrigued, concerned and befuddled researchers and policymakers for decades. Across the United States, among individuals 65 and over covered by a common health insurance program (Medicare), spending per capita across observably similar individuals varies by a factor of two; however, Medicare enrollees in higher-spending areas (such as Miami and McAllen) experience the same or worse health than in their lower-spending counterparts (such as Minneapolis and El Paso). In this paper, we ask whether places that have a larger *causal impact* on Medicare spending are also places that have a larger *causal impact* on enrollee’s life expectancy. We find that they do not, and discuss the implications.

The fact that places with higher spending per capita do not tend to experience better health outcomes has been widely interpreted as evidence of ‘flat-of-the-curve’ medicine – that there is little to no health return to higher medical spending, with the implication that health care spending could be substantially reduced in higher-spending parts of the country without harming pa-

tient health (e.g., Skinner, 2011; Office, 2008; Chandra, Cutler and Song, 2011; Gawande, 2009). This interpretation, in turn, informed the Obama administration’s view that they could achieve substantial cost savings – on the order of 20 to 30 percent of national health care spending or about 3 percent of GDP – without adverse effects on patient health, which in turn could pay for the health insurance expansions under the 2010 Affordable Care Act (e.g., Orszag, 2009; Pear, 2009).

This interpretation was premised on the idea that variation across areas in health care spending per capita primarily reflected differences in place-specific, supply-side variables, such as doctors’ practice styles or hospital market structure, rather than variation in patient characteristics such as their underlying health or preferences over treatments. If, in fact, patients in high-spending areas were simply sicker or preferred more intensive medical care, policies designed to reduce spending in those areas could be counterproductive. Therefore, a key issue was to isolate the role of place-based factors as opposed to person-based factors in driving the observed area variations in health care spending and in health.

To do so, we and our co-authors have exploited ‘mover designs’ that examine how outcomes change as individuals move across areas (Finkelstein, Gentzkow and Williams, 2016; Finkelstein et al., 2017; Finkelstein, Gentzkow and Williams, 2021; Finkelstein, Gentzkow and Li, 2025). To fix ideas, consider a Medicare enrollee who moves from a high-spending part of the country to a low-spending area. If all of the spending differences between those two places arise from place-based factors, we would expect the individual’s health care spending to drop immediately following the move, to a level similar to other patients in the low-

* Finkelstein: Massachusetts Institute of Technology, 50 Memorial Drive E52-442, Cambridge MA 02142, afink@mit.edu. Gentzkow: Stanford University, 579 Jane Stanford Way, Stanford CA 94305, gentzkow@stanford.edu. We are grateful to Peter Hull and Jonathan Skinner for helpful discussions, to Rosa Kleinman for excellent research assistance, and to the National Institute of Aging (Finkelstein, R01-AG032449) and the Peterson Center on Healthcare for financial support.

spending area. At the other extreme, if all of the spending differences between the two places reflect differences in the health and preferences of people living in the two areas, we would expect the individual's spending to remain constant after the move, at a level similar to the typical person in the high-spending area. Where the change in health care spending falls in between these two extremes can be used to identify the relative importance of person-based and place-based factors in driving regional variations. The same intuition can be applied in considering individuals who move from high-mortality parts of the country to low mortality parts (or vice versa).

Our findings indicated that place-based factors can explain about 50 to 60 percent of the observed geographic variation in Medicare spending per capita (Finkelstein, Gentzkow and Williams, 2016). We also found that place-based factors play an important role in affecting life expectancy at 65 (Finkelstein, Gentzkow and Williams, 2021), as did other work using a mover design (Deryugina and Molitor, 2021); for example, we estimated that relative to a 10th percentile place (in terms of its impact on life expectancy at 65), a 90th percentile place increased life expectancy at 65 by about 1.1 years, or about half of the 90-10 cross-sectional variation. However, we also found that place effects only explain about 15 percent of the geographic variation in life expectancy at 65, which is consistent with the conventional wisdom of a large role for person-based factors (particularly health behaviors) in driving much of the geographic variation in life expectancy (e.g., Fuchs, 1974; Cutler et al., 2011; Chetty et al., 2016).

Despite this progress in isolating place-specific causal effects on health care spending and on life expectancy, we do not yet know whether the well-documented lack of a relationship between area health care spending and area life expectancy changes when we move to examining the correlation between the place-specific causal components of health care spending and life expectancy. In other words, are places which have been estimated to have a higher causal

impact on health care spending per capita more likely to be places which have been estimated to have a higher causal impact on life expectancy? This paper rectifies that omission.

I. Empirical Approach

Our estimates of Medicare spending per capita and place effects on Medicare spending per capita are taken directly from Finkelstein, Gentzkow and Williams (2016). There, we used a 20 percent random sample of Medicare enrollees 65 and over from 1998 through 2008. As is typical in this literature, we restricted the analysis to Traditional Medicare enrollees, excluding the approximately one-fifth of Medicare enrollees enrolled in Medicare Advantage, a program in which private insurers receive capitated payments from the government and for whom we do not observe insurance claims. We partitioned the country into 306 Hospital Referral Regions (HRRs), which are collections of zip codes designed to approximate markets for hospital care. We used a measure of health care spending that is adjusted for area variation in administratively set prices, so that variation in health care spending reflects (price-weighted) variation in health care utilization. Our primary outcome was the log of price adjusted spending per capita. We estimated place effects on this outcome using a two-way fixed effects model of persons and places, with Medicare enrollees who move across places identifying the place-based and person-based drivers of spending.

Our measures of life expectancy at 65 and place effects on life expectancy at 65 are based on the approach in Finkelstein, Gentzkow and Williams (2021), with two important adjustments. First, we re-do our analyses at the HRR-level (rather than using Commuting Zones, as in that paper), so that spending and health are measured for the same area definitions. Second, that paper's estimates were based on a 100 percent sample of Medicare enrollees ages 65 and over from 1998 through 2014. Here, because we want to adjust both spending and life expectancy place effect estimates

for sampling error using a standard empirical Bayes approach, we exclude from our analysis of life expectancy individuals in the 20 percent random sample who were enrolled in Traditional Medicare, since these individuals who were used in estimating place effects in Finkelstein, Gentzkow and Williams (2016); this ensures that sampling errors for the two place-based estimates are independent.

Otherwise, our estimation strategy for life expectancy and place effects on life expectancy follows Finkelstein, Gentzkow and Williams (2021) exactly. To estimate average area life expectancy at 65, we estimate a Gompertz model of mortality in which the log of the individual's mortality rate is linear in age; we include an indicator for their area of residence, as well as controls for gender, race and their interaction, and indicator variables for the presence of about two dozen chronic conditions. To estimate place effects on life expectancy at 65, we leverage individuals who move across areas and compare how mortality varies across individuals who move to different locations; we control for their origin locations and a rich vector of pre-move health measures, and we adjust for any remaining selection on unobserved health using a novel strategy based on the correlation between movers' choice of destination and their observed health capital, as well as the correlation between residual post-move mortality with movers' origins to gauge the importance of omitted variables.

Finally, we adjust both sets of place-based estimates for sampling error using a standard empirical Bayes approach (see e.g. (Chetty and Hendren, 2018; Finkelstein et al., 2017); Appendix A of Finkelstein, Gentzkow and Williams (2021) provides more details.

II. Results

Figure 1 shows the relationship between our estimates of average area life expectancy at 65 across HRRs and average log Medicare spending per capita. It shows large variation across areas in both measures; however, areas with higher spending per capita tend to exhibit lower life ex-

pectancy.¹

Figure 2 shows the analogous relationship between our estimates of place effects on life expectancy and place effects on log Medicare spending. There is also substantial heterogeneity across places in their effect on health care spending (x-axis) and in their effect on life expectancy at 65 (y-axis). However, places that produce more health care spending per capita among the elderly do not produce higher life expectancy.

III. Discussion

Results like those in Figure 1 that demonstrate that areas with higher spending per capita did not experience better health outcomes have been extremely influential, but also controversial. While some interpreted this finding as evidence that health care spending – and hence taxpayer spending – could be reduced substantially without harming patients, others cautioned that the relationship might be contaminated by unobserved factors that are positively correlated with health care spending and negatively correlated with health.

The results in Figure 2 rule out one of the most natural potential confounders in Figure 1: unobserved differences in Medicare enrollees across areas. In particular, any differences across areas in the latent sickness or frailty of enrollees that could not be measured and controlled for would tend to drive up area health care spending and drive down area life expectancy. By correlating causal effects of place on health care spending with causal effects of place on health, the results in Figure 2 avoid this type of confounding. But they still admit two very different interpretations.

One interpretation of Figure 2 is that it validates the original interpretation of Figure 1 as indicative of ‘flat-of-the-curve’ medicine. The fact that areas that have a larger causal impact on health care spend-

¹If we instead examine the relationship between life expectancy at 65 and the average *level* of Medicare spending per capita, they are essentially uncorrelated. We show the relationship to average log Medicare spending here so as to be directly comparable with the place effects analysis below.

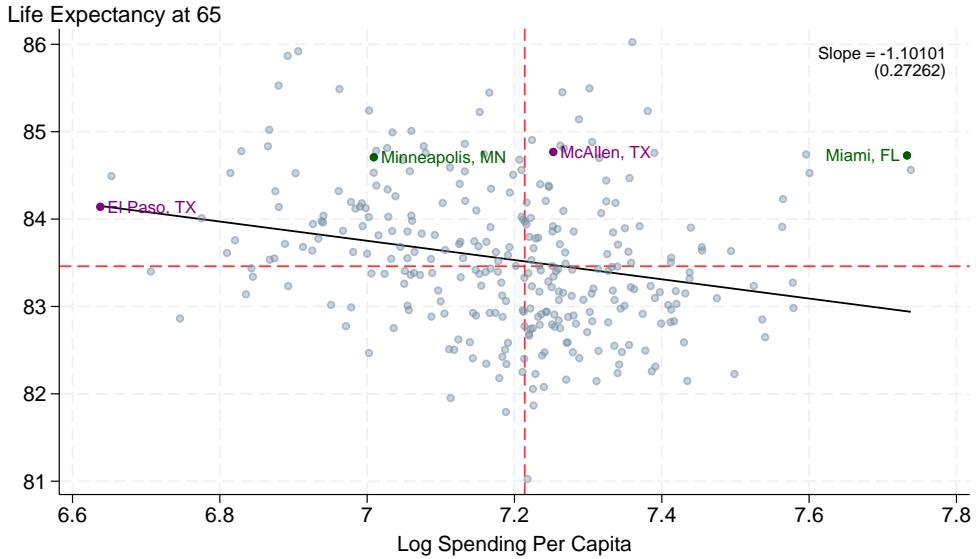


FIGURE 1. LIFE EXPECTANCY AT 65 AND ADJUSTED MEDICARE-SPENDING PER CAPITA

Note: Figure shows a scatterplot of HRR-average age 65 life expectancy (adjusted for race, sex and the presence of various health conditions) against HRR average log price-adjusted spending per capita (adjusted for age, race and sex). The line of best fit comes from a regression of life expectancy on log spending per capita, with slope and robust standard error (in parentheses) displayed in the top right corner. Dashed horizontal and vertical red lines indicate the medians of log spending per capita and life expectancy, respectively, across all HRRs.

ing do not have a larger causal impact on health is consistent with the idea that, on the margin, there is little or no ‘health return’ to higher medical spending. Conceptually, such a finding can be explained by excessive demand for health care from insured individuals who do not internalize the cost of that care (e.g. Feldstein (1973)) as well as by excess supply of health care from health care providers who benefit financially on the margin from more intensive medical treatment (e.g. Ellis and McGuire (1986)). Viewed through this lens, the key question is how to integrate this evidence of ‘flat-of-the-curve’ medicine with other studies using quasi-experimental variation in the intensity of medical care that have estimated positive health returns to increased medical care in specific contexts (e.g., Miller, Johnson and Wherry, 2021; Wyse and Meyer, 2025).

However, a very different interpretation of Figure 2 is that it mainly reflects place-specific factors that happen to be correlated with both a place’s impact on health

and its impact on spending. One possible confounder, for which there is growing evidence, is differences across areas in the efficiency or productivity of their health care system (Chandra et al., 2016; Chandra, Colla and Skinner, 2023). Assigning a patient to a high-productivity area could both improve their health and reduce their health spending, if the area is able to produce more output (health) for a given amount of input (spending). This could lead to a flat or even negative relationship in Figure 1, even if the marginal return to additional spending within any given area is positive. Consistent with this hypothesis, there is evidence that lower productivity clinical teams tend to engage in higher spending (Doyle, Ewer and Wagner, 2010).

Policy efforts under the Affordable Care Act aim to ‘bend the cost curve’ by reducing high-spending areas like Miami and McAllen closer to the levels of their lower-spending counterparts in Minneapolis and El Paso. Such efforts make sense under the ‘flat-of-the-curve’ interpretation of Fig-

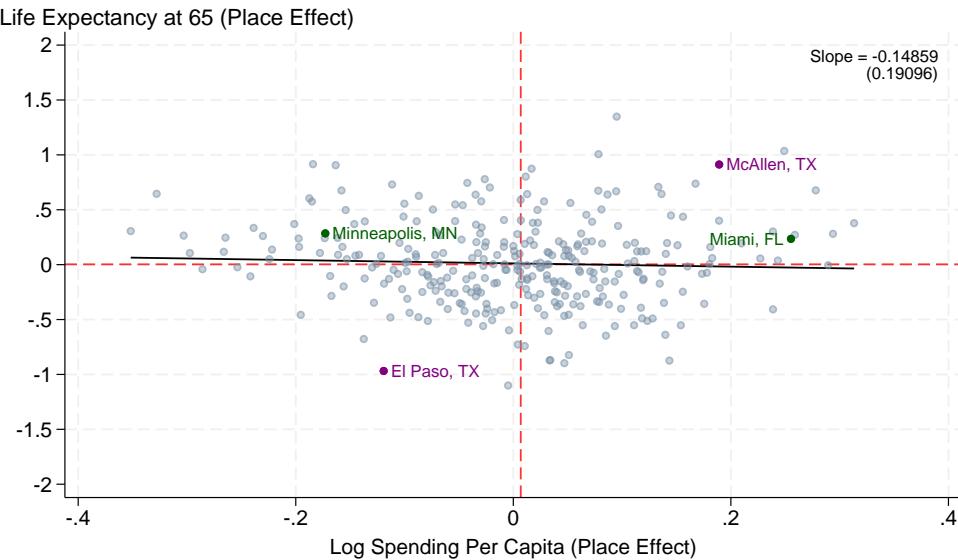


FIGURE 2. LIFE EXPECTANCY *Place Effects* AT 65 AND PER CAPITA SPENDING *Place Effects*

Note: Figure shows a scatterplot of life expectancy place effects against spending place effects, computed at the HRR level. Both measures are empirical-Bayes adjusted. The line of best fit comes from a regression of life expectancy place effects on spending place effects, with slope and robust standard error (in parentheses) displayed in the top right corner. Dashed horizontal and vertical red lines indicate the medians of spending and life expectancy place effects, respectively, across all HRRs.

ure 1. However, if other unobserved factors, such as health care productivity, are key drivers of the relationship between area health care spending and area health outcomes, it is possible that the marginal health return to spending in Miami could still be quite high and efforts to reduce spending could have negative consequences for patients. The results in Figure 2 indicate that differences in patient characteristics are not the key driver of the lack of a positive relationship between area spending and area health outcomes. Whether the 'flat-of-the-curve' interpretation is the correct one, however, remains an important area for further research.

REFERENCES

Chandra, Amitabh, Amy Finkelstein, Adam Sacarny, and Chad Syverson. 2016. "Health care exceptionalism? Performance and allocation in the US health care sector." *American Economic Review*, 106(8): 2110–2144.

Chandra, Amitabh, Carrie H Colla, and Jonathan S Skinner. 2023. "Productivity Variation and Input Misallocation: Evidence from Hospitals." National Bureau of Economic Research.

Chandra, Amitabh, David Cutler, and Zirui Song. 2011. "Who ordered that? The economics of treatment choices in medical care." *Handbook of health economics*, 2: 397–432.

Chetty, Raj, and Nathaniel Hendren. 2018. "The impacts of neighborhoods on intergenerational mobility II: County-level estimates." *The Quarterly Journal of Economics*, 133(3): 1163–1228.

Chetty, Raj, Michael Stepner, Sarah Abraham, Shelby Lin, Benjamin Scuderi, Nicholas Turner, Augustin Bergeron, and David Cutler. 2016. "The association between income and life expectancy in the United States, 2001–2014." *Jama*, 315(16): 1750–1766.

Cutler, David M., Fabian Lange, Ellen Meara, Seth Richards-Shubik, and Christopher J. Ruhm. 2011. "Rising Educational Gradients in Mortality: The Role of Behavioral Risk Factors." *Journal of Health Economics*, 30(6): 1174–87.

Deryugina, Tatyana, and David Moller. 2021. "The Causal Effects of Place on Health and Longevity."

Doyle, Joseph J, Steven M Ewer, and Todd H Wagner. 2010. "Returns to physician human capital: Evidence from patients randomized to physician teams." *Journal of health economics*, 29(6): 866–882.

Ellis, Randall P, and Thomas G McGuire. 1986. "Provider behavior under prospective reimbursement: Cost sharing and supply." *Journal of health economics*, 5(2): 129–151.

Feldstein, Martin S. 1973. "The welfare loss of excess health insurance." *Journal of Political Economy*, 81(2, Part 1): 251–280.

Finkelstein, Amy, Matthew Gentzkow, and Dean Li. 2025. "What drives risky prescription opioid use? Evidence from migration." *The quarterly journal of economics*, 140(4): 3133–3189.

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

Finkelstein, Amy, Matthew Gentzkow, and Heidi L. Williams. 2021. "Place-Based Drivers of Mortality: Evidence from Migration." *American Economic Review*, 111(8): 2697–2735.

Finkelstein, Amy, Matthew Gentzkow, Peter Hull, and Heidi Williams. 2017. "Adjusting risk adjustment—accounting for variation in diagnostic intensity." *The New England journal of medicine*, 376(7): 608.

Fisher, Elliott S, David E Wennberg, Thrse A Stukel, Daniel J Gottlieb, F Lee Lucas, and Etoile L Pinder. 2003a. "The implications of regional variations in Medicare spending. Part 1: the content, quality, and accessibility of care." *Annals of internal medicine*, 138(4): 273–287.

Fisher, Elliott S, David E Wennberg, Thrse A Stukel, Daniel J Gottlieb, F Lee Lucas, and Etoile L Pinder. 2003b. "The implications of regional variations in Medicare spending. Part 2: health outcomes and satisfaction with care." *Annals of internal medicine*, 138(4): 288–298.

Fuchs, Victor R. 1974. *Who Shall Live? Health, Economics and Social Choice*. New York:Basic Books.

Gawande, Atul. 2009. "The Cost Conundrum." *The New Yorker*.

Miller, Sarah, Norman Johnson, and Laura R. Wherry. 2021. "Medicaid and Mortality: New Evidence from Linked Survey and Administrative Data." *Quarterly Journal of Economics*, 136(3): 1783–1829.

Office, Congressional Budget. 2008. "Geographic variation in health care spending."

Orszag, Peter R. 2009. "Testimony of Peter R. Orszag, Director of the Office of Management and Budget Before the Committee on Finance, U.S. Senate." United States Senate.

Pear, Robert. 2009. "Health Care Spending Disparities Stir a Fight." *The New York Times*.

Skinner, Jonathan. 2011. "Causes and consequences of regional variations in health care." In *Handbook of health economics*. Vol. 2, 45–93. Elsevier.

Wyse, Angela, and Bruce D Meyer. 2025. "Saved by Medicaid: new evidence on health insurance and mortality from the universe of low-income adults." National Bureau of Economic Research.