Medicare Reimbursement and Hospital Scale

Hieu M. Nguyen

Department of Economics and Business, Southwestern University

Abstract

Medicare routinely adjusts hospital reimbursement for local labor costs through the hospital wage index (HWI). I study whether discontinuities in this payment adjustment are associated with hospital scale. I combine administrative wage-index files, hospital cost reports, hospital characteristics, Census county geometry, and hospital locations, and employ a spatial regression discontinuity design to compare hospitals near the same wage-index-area boundary. I find evidence consistent with a scale relationship. Hospitals on the higher wage-index side are larger within narrow local distances, especially in FTE staff, beds, payroll, and total costs, yet the estimates attenuate in wider windows.

Keywords: Medicare; hospital wage index; hospital employment; spatial regression discontinuity; geographic payment adjustment

JEL Codes: I11, I18, J20

List of abbreviations: CBSA, Core Based Statistical Area; CMS, Centers for Medicare and Medicaid Services; DRG, Diagnosis-Related Group; FTE, Full-Time Equivalent; FY, Fiscal Year; HIFLD, Homeland Infrastructure Foundation-Level Data; HWI, Hospital Wage Index; ICPSR, Inter-university Consortium for Political and Social Research; IPPS, Inpatient Prospective Payment System; OMB, Office of Management and Budget; PPS, Prospective Payment System; RDD, Regression Discontinuity Design.

Conflict of interest: None.

Funding statement: None.

Correspondence: Hieu Nguyen, 1001 E University Ave, Georgetown, TX 78626. Email: nguyen23@southwestern.edu.

1. Introduction

Medicare's Inpatient Prospective Payment System (IPPS) adjusts hospital payments for local labor costs through what is known as the "hospital wage index" (HWI). The HWI has been a feature of IPPS since the system's introduction under the Social Security Amendments of 1983, which replaced *cost-based* reimbursement with *per-discharge* payments based on diagnosis-related groups (DRGs) starting in October 1983. The wage index is applied to the labor-related share of the base payment rate, which has historically accounted for approximately 60 to 70 percent of the standardized amount. It varies across wage index areas, which correspond to Core Based Statistical Areas (CBSAs) for urban hospitals and statewide rural designations for hospitals outside metropolitan areas.

From a conceptual standpoint, hospital location and operational efficiency are not randomly assigned, but are likely endogenous to differences in labor costs, urbanization, patient demand, and, most importantly, the same factors that determine the wage index itself. An unconditional comparison between hospitals with a high wage index and those with a low wage index would therefore be potentially misleading. I instead employ a spatial regression discontinuity design (RDD) to compare hospitals physically located close to the same wage-index-area boundary. A hospital located just inside one wage index area can face a different payment adjustment from a geographically close hospital on the other side of the boundary. I ask whether these local payment differences are associated with hospital scale.

This analysis warrants attention for two reasons. First, labor is an important ingredient of the hospital production function. If geographic payment rules are correlated with hospital employment, capacity constraints, and other hospital-related outcomes, then Medicare payment design may have important implications for local healthcare provision beyond reimbursement accounting. Second, geographic adjustment is widely used in public payment systems, and its local consequences are difficult to measure because of the conceptual challenges outlined above. To my knowledge, this is the first study to exploit the geographic discontinuity in Medicare HWI values at area boundaries to study hospital-level outcomes.

This paper contributes to three related strands of work. First, it contributes to research on Medicare payment design and hospital behavior. A substantial literature documents

that payment rules shape hospital choices. Acemoglu and Finkelstein (2008) show that the replacement of full cost reimbursement with partial cost reimbursement under PPS significantly raised capital-labor ratios and lowered labor inputs, with hospitals substituting technology for staff following the regulatory change. Dafny (2005) shows that hospitals responded to diagnosis-specific Medicare price changes primarily by upcoding patients into codes with larger reimbursement increases. Lindrooth et al. (2006) show that Medicare reimbursement pressure can affect nurse staffing decisions. Kaestner and Guardado (2008) use geographic reclassification to study Medicare reimbursement changes, finding no meaningful effects on hospital resource use or patient outcomes. Cook and Averett (2020) study Medicare's 2007 DRG restructuring and find that hospitals responded to the new incentive structure through upcoding, including spillovers beyond Medicare patients.[1] Gross et al. (2024) study a related DRG revision and find that hospitals facing payment gains expanded nurse employment and patient volume. Relatedly, Clemens and Gottlieb (2014) show that Medicare physician payment shocks affect care provision, suggesting that public reimbursement formulas can change provider behavior. HWI-focused research has shown that area-level payment adjustments are associated with nursing wages and nursing home staffing (He et al., 2020; McHenry and Mellor, 2018). The identification strategies of these papers rely mostly on price changes, reclassification events, or area-wide variation, each of which involves comparisons across broad regions. By comparing hospitals close to the same wage-index-area boundary within a distance-based RDD framework, I therefore contribute an internal validity improvement to this strand of literature. Unlike the studies on upcoding, my paper does *not* study whether hospitals change reported patient severity in response to reimbursement incentives. I instead study whether a geographic payment formula is associated with the size of hospital operations.

Second, this paper expands what we know about health care organization and hospital labor markets. To my knowledge, no studies have looked at the potential nexus between the HWI and hospital scale. At its core, the question of hospital scale, concerning whether higher payment generosity is associated with larger hospitals in terms of beds, employment, and total costs, fundamentally differs from wage pass-

[1] The "upcoding" literature includes major papers such as Silverman and Skinner (2004), Jürges and Köberlein (2015), Shigeoka and Fushimi (2014), and Geruso and Layton (2020), among others. Cook and Averett (2020) briefly review this work in the Medicare DRG context.

through, and has received less direct attention in the HWI context. The Institute of Medicine (2012) notes that the HWI is central to Medicare's geographic adjustment of fee-for-service payments, but its relationship to the level of hospital operations at local boundaries has not been fully examined.

Third, this paper contributes to extending the empirical evidence on spatially informed economic evaluations. The approach of comparing units on opposite sides of an administrative boundary dates at least to Black (1999), who uses school attendance boundaries to study parental valuation of school quality, and Holmes (1998), who documents sharp differences in manufacturing activity at right-to-work state borders. Dell (2010) applies a spatial regression discontinuity to study the persistent effects of Peru's colonial mining institution, and Schumann (2014) uses the German occupation zone boundary to estimate the persistence of population shocks. I apply this class of designs to the HWI area boundary, which creates a well-defined policy discontinuity within otherwise continuous economic geography. The value of this design is that, by restricting comparisons to nearby hospitals, it holds fixed much of the broad regional environment that would otherwise confound unconditional comparisons.

This paper is closest in spirit to Kaestner and Guardado (2008) and McHenry and Mellor (2018). Like these two papers, I treat the HWI as a significant input into provider resource decisions. However, my investigation differs in two key respects. First, both these papers identify payment variation from reclassification events or area-level wage index variation, allowing them to study reimbursement changes across hospitals or labor markets. By design, these analyses compare hospitals across potentially heterogeneous geographies. The boundary design I employ mitigates the potential influence of spatially correlated unobserved confounders. Second, the prior literature finds some transmission from the HWI to worker pay. I do not replicate that pattern here. Rather, my results indicate stronger support for the impacts on the size of hospital operations. This shows one way in which Medicare payment geography may potentially affect local health care capacity through channels unrelated to wages.

2. Institutional Context

The IPPS was enacted through the Social Security Amendments of 1983 and has since governed payment for most acute care hospital inpatient stays. The system replaced cost-based reimbursement with per-discharge payments based on diagnosis-related

groups. For each Medicare discharge, a hospital receives a base payment that is adjusted by the DRG weight, the geographic wage index, and other factors such as indirect medical education and disproportionate share adjustments. The HWI applies to the labor-related share of the standardized amount, which CMS updates annually in its final rule.

The HWI is calculated from hospital cost reports and the Occupational Mix Survey. For each wage index area, CMS computes an area average hourly wage using reported wages and hours from hospitals in that area. CMS then adjusts for occupational mix to account for differences in the skill composition of nursing labor and divides the result by the national average hourly wage. An HWI above 1.0 indicates that the area's hospital labor costs are above the national average. An HWI below 1.0 indicates the reverse. CMS also applies several additional wage-index policies, including the rural floor, out-migration adjustment, MGCRB reclassification, and the permanent cap on annual wage-index decreases. These provisions affect payment, but are distinct from the baseline county-to-area assignment that motivates my subsequent empirical design.

County membership determines a hospital's baseline wage index area. Metropolitan counties are assigned to the CBSA in which they fall, whereas non-metropolitan counties are assigned to the statewide rural area for their state. CBSA definitions are established by OMB based on county-level commuting patterns and economic integration, not by CMS and not by hospitals. The boundaries are updated periodically by OMB and adopted by CMS for payment purposes. A major realignment occurred in fiscal year 2005, when CMS moved from earlier area definitions to CBSA-based definitions. McHenry and Mellor (2018) and He et al. (2020) use this realignment for identification. I instead use the FY 2026 wage index, which reflects current wage-index data and current area assignments. It is also worth noting that hospitals may petition for geographic reclassification through the Medicare Geographic Classification Review Board. If approved, a reclassified hospital is still located in its original county, but paid using the wage index of another area. Reclassification is therefore administratively distinct from the county assignment that defines the boundary itself.[2]

[2] In the analysis, I use the provider-level wage index because it is the most relevant payment parameter. However, I do acknowledge that *hospital-specific* reclassification can weaken the connection between physical location and payment assignment.

3. Data and Identification Strategy

3.1. Data

I combine six data sources for the empirical analysis. The most important source is the CMS FY 2026 Final Rule Wage Index Public Use File, which provides provider-level HWI values and the area-level wage information used to construct wage-index-area measures. To map each county to its wage index area, I use the CMS FY 2026 Final Rule county-to-CBSA crosswalk, which assigns counties to CBSAs for metropolitan areas and to statewide rural areas for non-metropolitan counties. The third data source is the CMS Hospital Cost Report Public Use File for fiscal year 2022, which supplies hospital employment, beds, discharges, payroll, costs, charges, and wage-related cost data. For information on hospital type and ownership, I use the CMS Hospital General Information file. To approximate wage index area boundaries, I use the Census 2024 TIGER/Line county shapefile, which contains county boundary coordinates. Finally, I use the restricted-access HIFLD Open Hospitals shapefile made available by DataLumos. [3]

To obtain the final analytic dataset, I proceed in several stages. First, I merge the wage index provider file to the cost report file using the CMS Certification Number. I then assign each hospital to a wage index area using the county-to-CBSA crosswalk, with non-CBSA counties classified as statewide rural areas following CMS convention. Next, I link hospital coordinates from the HIFLD DataLumos file by matching on state, county FIPS code, and cleaned hospital name. The cleaning step removes common legal suffixes and generic hospital descriptors to improve match rates while mitigating inaccurate matches.[4] In the last stage, I compute signed distance to the nearest different wage index area. A practical challenge here is that the treatment changes across wage-index-area geography, while hospitals are observed as points. I therefore first identify, for each hospital, the nearest county centroid assigned to a different wage index area. I then use the county shapefile to recover the vertices of that comparison county and

[3] This file is accessed through the ICPSR DataLumos portal. This particular file is contained within a zip folder that requires researcher registration and is not directly downloadable without a formal application.
[4] The current public files do not allow me to handle hospitals with multiple campuses with complete certainty, since a CMS Certification Number (CCN) can represent a provider organization, while HIFLD records physical facilities. I therefore use a conservative coordinate-matching procedure. Specifically, the matching algorithm requires agreement on state, county FIPS code, and cleaned hospital name. This requirement helps reduce inaccurate matches, although it may exclude some valid multi-campus providers.

calculate the distance from the hospital point to the nearest vertex using the Haversine formula. The resulting distance is signed according to whether the hospital's own wage index is above or below the nearest comparison area's wage index.[5] This two-step procedure uses centroids to identify the relevant neighboring area and vertices to approximate boundary distance more precisely than centroid-to-centroid calculations allow. It is still an approximation to the true polygon boundary distance, but it is a refinement that moves the running variable closer to the spatial margin implied by the administrative payment rule. This entire process produces a conservative matched sample of 1,097 hospitals.

3.2. Identification Strategy

Let i index hospitals. Let d_i denote signed distance in kilometers from hospital i to the nearest county boundary with a different wage index value. Positive values indicate the higher wage index side. Negative values indicate the lower side. Let H_i be a binary indicator equaling one for hospitals on the higher wage index side and zero otherwise.

The main estimating equation is:

$$Y_i = \alpha + \tau H_i + \beta d_i + \gamma H_i \times d_i + \epsilon_i \qquad (1),$$

where Y_i is a winsorized untransformed outcome for hospital i. My main outcome of interest is full-time equivalent (FTE) employment, although I will also examine beds, discharges, total salaries, total costs, total charges, and wage cost share.[6] The coefficient τ is the discontinuity at the boundary. The term d_i allows a linear distance slope on the low side. The interaction term allows that slope to differ on the high side. I also estimate an alternative specification that adds the winsorized number of licensed beds to the right-hand side of equation (1), which allows me to explore whether the potential

---

[5] The Haversine distances are calculated as follows. For a hospital point i and a comparison point j, with latitudes $\phi_i; \phi_j$ and longitudes $\lambda_i; \lambda_j$ measured in radians, I define $a_{ij} = \sin^2\left(\frac{\phi_j - \phi_i}{2}\right) + \cos(\phi_i)\cos(\phi_j)\sin^2\left(\frac{\lambda_j - \lambda_i}{2}\right)$, and great-circle distance $d_{ij} = 2R\sin^{-1}\left(\sqrt{a_{ij}}\right)$, where R represents the Earth's radius. In Stata, I use $R = 3958.7613$ miles and convert the final distance to kilometers. The running variable keeps the minimum distance to the relevant comparison county vertex and signs it by the wage-index gap.

[6] I winsorize all outcomes at the first and 99th percentiles and use levels rather than logarithms to make the estimates easily interpretable in natural units. Salary per FTE is not a main outcome because the wage per worker specifications do not produce a consistent pattern (comprising mostly null results) across distance windows. The results for this exercise are available upon request.

employment, payroll, and cost differences persist after conditioning on hospital capacity. For robustness and sensitivity checks, I estimate equation (1) within windows of 5, 10, 15, 20, and 25 kilometers. I also report standard RDD diagnostics in the appendix. Appendix A2 plots the density of the signed-distance running variable within 25 kilometers of the boundary. The graph shows a smooth transition through zero with no sharp bunching at the cutoff.[7] Since the running variable is physical distance to an administrative boundary and not a hospital-chosen score, this diagnostic is not a direct test of manipulation in the conventional sense. Nevertheless, the result confirms that hospitals are *not* more concentrated on one side of the boundary than the other within the local window, and helps mitigate the concern that hospitals might have endogenously located on the more generously reimbursed side in anticipation of payment differences. Appendix A3 reports falsification checks using observable ownership categories available in the CMS Hospital General Information file. The table shows no evidence of discontinuities in nonprofit, proprietary, or government ownership at the boundary, providing some reassurance that the local comparison is not driven by sharp ownership-related sorting.

Several other considerations are worth discussing. First, wage index area boundaries are not randomly assigned. They capture labor market geography, and hospitals near boundaries may differ because of local economic conditions rather than Medicare payment. My restriction to comparisons involving narrow distance windows of 5 to 25 kilometers and my examination of licensed beds as a scale check should partially alleviate this concern. Relatedly, some wage-index-area boundaries coincide with state lines because wage-index areas are built from county assignments and because non-metropolitan hospitals are assigned to statewide rural areas. Comparisons across state-line boundaries therefore risk confounding the wage-index discontinuity with state-level policy differences that also change at the same border, including Medicaid reimbursement rates, state income tax regimes, and so forth. To address this, I construct a same-state restriction that retains only hospitals whose nearest comparison county belongs to the same state, thereby excluding comparisons in which the nearest different wage-index county is across a state line. Appendix A4 replicates Table 2 under this

[7] A formal test using the Stata command "rddensity" is also consistent with this visual pattern, with a test statistic of -0.132 ($p = 0.895$). This shows no evidence of a discontinuity in the density of hospitals at the wage-index-area boundary. The 295 hospitals on the lower wage-index side and 283 hospitals on the higher wage-index side are almost symmetrically distributed around the normalized cutoff.

restriction. The estimates are similar in magnitude and statistical significance to the main results, and in several cases slightly larger, which is inconsistent with state-level policy heterogeneity driving the findings.

Second, from the institutional description one might be concerned that hospitals and populations may sort across boundaries. This concern cannot be ruled out completely, but it is arguably less severe in the current context because county boundaries and CBSA definitions are not chosen by hospitals. Hospitals also cannot continuously manipulate their distance to a boundary. More precisely, a hospital's wage-index area is determined entirely by the county in which it is physically located, and county membership is itself a fixed administrative fact derived from the hospital's address. A hospital that wishes to receive the wage index of an adjacent area cannot do so simply by relocating a small distance. It would need to cross a county line, which in practice means a discrete relocation of the facility itself. The reclassification process offers an administrative workaround, but as discussed in Section 2, reclassification affects only the payment index and not the hospital's physical location, so it does not change the running variable in my current design.

Finally, one might suspect that, given the non-score-based assignment mechanism, the associated estimates are difficult to interpret. However, a hospital is located in one county or another, and the payment discontinuity occurs at a known administrative line. The design is therefore closest to a sharp spatial RDD, with the treatment indicator being defined relative to a known administrative boundary and the running variable measured as signed physical distance to the nearest different wage-index area. The county-line assignment mechanism implies that the treatment is a step function of physical geography determined by an administrative rule that predates current payment levels and that no individual hospital can influence at the margin. This supports a local boundary interpretation.

4. Results

Table 1 displays the descriptive statistics for the analytic sample of 1,097 matched hospitals. The average wage index is 0.993, with a standard deviation of 0.186. Figure 1 shows that the distribution is concentrated around 1.0, but with meaningful variability across hospitals. The average wage-index gap between a hospital and its nearest different wage-index area is 0.006, with a standard deviation of 0.090. Figure 2 shows that most local gaps are close to zero, although there is meaningful variation on both

sides. This pattern is useful for the empirical design because the comparison is not driven only by a small number of extreme payment differences. The average hospital in the matched sample has 1,631 FTE employees, 235 beds, and 10,955 annual discharges. Average salaries are $152 million, and average total costs are $350 million.

Table 2 reports the local boundary estimates from equation (1). The clearest result appears in the narrowest window. Within 5 kilometers of a wage-index-area boundary, hospitals on the higher wage-index side have approximately 3,674 more FTE employees than hospitals on the lower wage-index side. The estimate is statistically significant at the 5 percent level of statistical significance. The corresponding 10-kilometer estimate is 1,756 FTE employees and is significant at the 10 percent level. The estimate then becomes smaller and is no longer statistically significant at 15, 20, or 25 kilometers. Figure 3 is consistent with this pattern. It shows a visible difference in binned mean FTE near the boundary, but also makes clear that the comparison is local and, as cautioned above, should not be interpreted as a broad comparison between all high and low wage-index hospitals. Appendix Figure A1 repeats the same graphical exercise using data-driven RD plots following Calonico et al. (2015). This figure is intended as a graphical robustness check, not as a separate estimating equation.

Similar patterns emerge for several scale-related outcomes. Within 5 kilometers, hospitals on the higher wage-index side have approximately 277 more beds, 12,301 more annual discharges, $392 million more in salaries, $754 million more in total costs, and $3.56 billion more in charges. These estimates are statistically significant at the 5 percent level. At 10 kilometers, the salary and total cost estimates remain significant at the 10 percent level, but the other scale outcomes weaken. By 15 kilometers and beyond, most estimates are smaller and no longer statistically distinguishable from zero. The wage-cost-share estimates are small and statistically insignificant across all windows. This null result is somewhat informative, because it rules out a simplistic interpretation in which the higher wage index side merely spends a larger share of total costs on labor.

Table 3 asks whether these differences persist after conditioning on hospital capacity. This specification adds winsorized beds as a linear covariate. It is informative but should be interpreted as a diagnostic specification rather than as a second causal specification, since beds are themselves part of hospital scale. The 5 kilometer FTE estimate falls from 3,674 in Table 2 to 587 in Table 3 and is no longer statistically significant. At 10 kilometers, however, the bed-adjusted FTE estimate is 843 and

statistically significant at the 5 percent level. The estimates at 15 and 20 kilometers are smaller but remain significant at the 10 percent level.

The bed-adjusted cost results point in a similar direction. Total costs remain higher on the higher wage-index side at 10, 15, and 20 kilometers, with estimates of $162 million, $118 million, and $101 million, respectively. These estimates are statistically significant at the 5 percent level. Salaries are also higher at 10 kilometers, with an estimate of $87 million. In contrast, the bed-adjusted discharge estimates are negative and statistically insignificant across all windows, and wage-cost share remains statistically insignificant.

Together, these results are most consistent with a scale interpretation. In the smallest local window, hospitals on the higher wage-index side are larger in beds, FTE employment, discharges, salaries, total costs, and charges. Once I condition on beds, the narrow-window FTE difference largely disappears, which suggests that the main margin is scale as opposed to staffing intensity. At the same time, some cost differences persist after the bed adjustment, especially in wider distance windows. This suggests that Medicare payment geography is associated with hospital scale and cost levels near wage-index-area boundaries, and that a simple wage pass-through mechanism is unlikely to fully explain the results.

To help interpret the economic magnitude of the boundary estimates, I conduct a back-of-the-envelope calculation under the assumption of linearity. To be clear, this exercise is not a structural estimate or a welfare calculation. It simply *rescales* the reduced-form estimates by the mean absolute wage-index gap observed within each bandwidth. While this exercise neither allows me to predict the system-wide effect of a policy reform nor accounts for how hospitals in different markets might respond differently, it is useful in placing the magnitudes associated with wage-index differences in perspective. Specifically, for outcome S and bandwidth b, I calculate:

$$\widehat{\Delta S}_{0.01,b} = \hat{\tau}_{S,b} \times \left(\frac{0.01}{\overline{|g|}_b} \right) \qquad (2).$$

Here, $\hat{\tau}_{S,b}$ is the estimated discontinuity from Table 2, and $\overline{|g|}_b$ is the average absolute wage-index gap among hospitals within the same bandwidth. The quantity $\widehat{\Delta S}_{0.01,b}$ is therefore the implied difference in outcome S associated with a one-percentage-point wage-index gap.

Appendix A5 shows that the implied magnitudes are economically meaningful but highly local. In the 5 km window, a 0.01 wage-index gap corresponds to approximately 482 FTE staff, 36 beds, 1,612 annual discharges, $51 million in salaries, $99 million in total costs, and $466 million in charges. In the 10 km window, the implied magnitudes fall to 267 FTE staff, 13 beds, 541 annual discharges, $28 million in salaries, $53 million in total costs, and $211 million in charges. Consistent with the main reduced-form pattern, the implied differences become smaller in wider windows. These figures should not be interpreted as structural estimates, but the magnitudes are economically significant.

5. Concluding Remarks

Medicare's hospital wage index is designed to adjust payments for local labor costs. In this study, I have examined whether wage index area boundaries are associated with hospital scale using a spatial regression discontinuity design. The main pattern is that hospitals on the higher wage index side appear larger in the narrowest local windows. FTE employment is significantly higher at 5 and 10 kilometers, beds are higher at 5 kilometers, and cost differences persist after controlling for beds. Wage cost share shows no discontinuity at any window.

Several limitations are worth emphasizing. First, it is important to reemphasize that wage index area boundaries are not randomly assigned. They follow county assignments to CBSAs, which likely are endogenous to labor market structure and historical geography. Narrow distance windows do help reduce broad spatial confounding, but they cannot eliminate every preexisting local difference. The null result on wage cost share provides partial comfort that the estimates do not simply pick up local differences in labor market conditions. The overall evidence is therefore most interpretable as documenting equilibrium differences in hospital scale near boundaries, not as identifying the causal effect of a marginal wage index change. Longitudinal data combined with discrete changes in the wage index, such as the 2005 CBSA realignment studied by McHenry and Mellor (2018), would allow stronger causal identification than the cross-sectional boundary design considered here.

Second, the hospital location matching exercise relies on linking CMS administrative records to the HIFLD DataLumos coordinates by state, county FIPS code, and cleaned hospital name. This conservative match yields 1,097 hospitals out of over 7,000 total.

The exclusion of unmatched hospitals could introduce selection bias if match rates differ systematically by wage index level or distance to the boundary. The advantage of using actual hospital coordinates is that it provides a more precise distance measure than county centroid approximations would provide. Having said that, this limitation has a less direct effect on internal validity, so long as matching error is not discontinuous at the wage-index boundary.

Third, the analysis cannot distinguish whether the scale differences reflect active responses to the current payment environment or persistent historical patterns that predate the current wage index values. This limitation is shared with other canonical spatial RDDs; however, it does not make the exercise less informative. The results identify a unique dimension of variation at HWI boundaries that has not previously been studied. This last limitation should motivate future work, especially studies that leverage longitudinal designs to track how hospital scale responds to changes in the wage index over time.

References

Acemoglu, D., & Finkelstein, A. (2008). Input and technology choices in regulated industries: Evidence from the health care sector. *Journal of Political Economy*, *116*(5), 837-880.

Black, S. E. (1999). Do better schools matter? Parental valuation of elementary education. *The Quarterly Journal of Economics*, *114*(2), 577-599.

Calonico, S., Cattaneo, M. D., & Titiunik, R. (2015). Optimal data-driven regression discontinuity plots. *Journal of the American Statistical Association*, *110*(512), 1753-1769.

Centers for Medicare & Medicaid Services. (2025). Medicare Program; Hospital Inpatient Prospective Payment Systems for Acute Care Hospitals, IPPS and LTCH PPS policy changes and fiscal year 2026 rates; changes to the FY 2025 IPPS rates due to court decision; requirements for quality programs; and other policy changes. *Federal Register,* 90, 36536. Retrieved from: https://www.federalregister.gov/d/2025-14681.

Clemens, J., & Gottlieb, J. D. (2014). Do physicians' financial incentives affect medical treatment and patient health?. *American Economic Review*, *104*(4), 1320-1349.

Cook, A., & Averett, S. (2020). Do hospitals respond to changing incentive structures? Evidence from Medicare's 2007 DRG restructuring. *Journal of Health Economics*, *73*, 102319.

Dafny, L. S. (2005). How do hospitals respond to price changes?. *American Economic Review*, *95*(5), 1525-1547.

Dell, M. (2010). The persistent effects of Peru's mining mita. *Econometrica*, *78*(6), 1863-1903.

Geruso, M., & Layton, T. (2020). Upcoding: evidence from Medicare on squishy risk adjustment. *Journal of Political Economy*, *128*(3), 984-1026.

Gross, T., Sacarny, A., Shi, M., & Silver, D. (2024). Regulated revenues and hospital behavior: Evidence from a medicare overhaul. *Review of Economics and Statistics*, *106*(6), 1709-1718.

He, D., McHenry, P., & Mellor, J. M. (2020). The effects of Medicare payment changes on nursing home staffing. *American Journal of Health Economics*, *6*(4), 411-443.

Holmes, T. J. (1998). The effect of state policies on the location of manufacturing: Evidence from state borders. *Journal of Political Economy*, *106*(4), 667-705.

Institute of Medicine. (2012). *Geographic Adjustment in Medicare Payment: Phase I: Improving Accuracy*. National Academies Press. Retrieved from: https://doi.org/10.17226/13138.

Jürges, H., & Köberlein, J. (2015). What explains DRG upcoding in neonatology? The roles of financial incentives and infant health. *Journal of Health Economics*, *43*, 13-26.

Kaestner, R., & Guardado, J. (2008). Medicare reimbursement, nurse staffing, and patient outcomes. *Journal of Health Economics*, *27*(2), 339-361.

Lindrooth, R. C., Bazzoli, G. J., Needleman, J., & Hasnain-Wynia, R. (2006). The effect of changes in hospital reimbursement on nurse staffing decisions at safety net and nonsafety net hospitals. *Health Services Research*, *41*(3p1), 701-720.

McHenry, P., & Mellor, J. (2018). Medicare hospital payment adjustments and nursing wages. *International Journal of Health Economics and Management*, *18*(2), 169-196.

Schumann, A. (2014). Persistence of population shocks: Evidence from the occupation of West Germany after World War II. *American Economic Journal: Applied Economics*, *6*(3), 189-205.

Shigeoka, H., & Fushimi, K. (2014). Supplier-induced demand for newborn treatment: evidence from Japan. *Journal of Health Economics*, *35*, 162-178.

Silverman, E., & Skinner, J. (2004). Medicare upcoding and hospital ownership. *Journal of Health Economics*, *23*(2), 369-389.

Table 1. Descriptive Statistics for the Analytic Sample

Variable	Mean	SD
Wage index	0.993	0.186
Wage-index gap	0.006	0.090
Signed distance to nearest different wage-index area, km	4.144	43.409
FTE employment	1,631.135	2,023.969
Beds	235.349	206.667
Annual discharges	10,954.536	10,293.020
Salaries, $ millions	152.441	205.599
Total costs, $ millions	349.726	424.546
Charges, $ millions	1,729.460	2,061.451
Wage-cost share	0.089	0.029

Notes. The table reports summary statistics for the 1,097 matched hospitals in the boundary analysis sample. The wage index is the hospital's assigned FY 2026 CMS wage index value. The wage-index gap is the difference between the hospital's own wage index and that of the nearest county with a different wage-index value. Signed distance is measured in kilometers from the hospital to the nearest boundary vertex of the nearest different wage-index county, with positive values indicating the higher wage-index side. Outcomes are winsorized at the 1st and 99th percentiles. Salaries, total costs, and charges are reported in millions of dollars.

Table 2. Local Boundary Estimates

Outcome	5 km	10 km	15 km	20 km	25 km
FTE employment	3,673.990**	1,756.152*	838.425	733.796	403.760
	(1,596.306)	(947.946)	(696.138)	(546.180)	(472.151)
Beds	276.615**	86.103	30.476	30.054	27.221
	(124.368)	(75.321)	(56.604)	(45.439)	(39.245)
Annual discharges	12,300.627**	3,556.806	657.196	608.071	740.546
	(6,116.681)	(3,719.274)	(2,753.393)	(2,241.212)	(1,949.047)
Salaries, $ millions	391.511**	183.935*	75.628	74.110	54.411
	(167.267)	(99.736)	(73.890)	(58.172)	(49.438)
Total costs, $ millions	753.918**	350.338*	180.266	160.290	114.817
	(326.003)	(192.646)	(141.819)	(111.015)	(94.507)
Charges, $ millions	3,555.265**	1,386.889	462.085	401.594	476.680
	(1,568.621)	(904.072)	(682.604)	(542.798)	(462.700)
Wage-cost share	0.003	0.006	0.008	0.010	0.009
	(0.015)	(0.009)	(0.008)	(0.007)	(0.006)
N	76	202	333	463	578

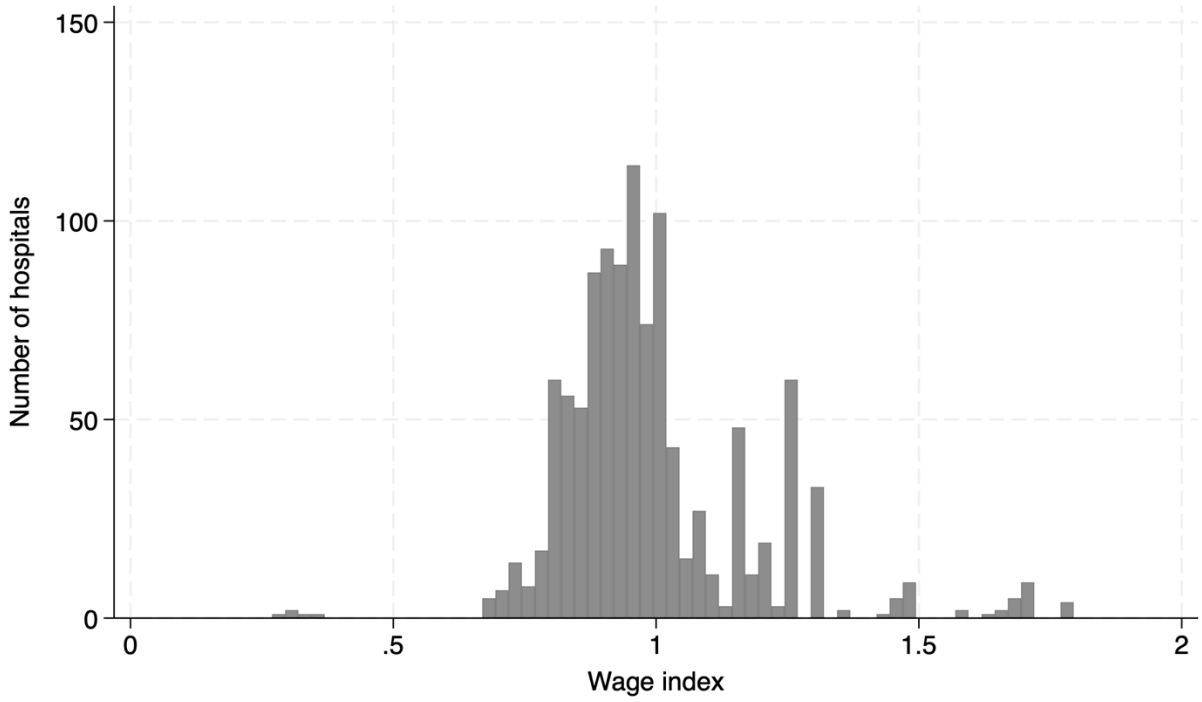
Notes. Each cell reports the coefficient on the higher wage-index side from a local linear boundary regression at the indicated bandwidth. The running variable is signed distance in kilometers to the nearest different wage-index area, with positive values indicating the higher wage-index side. Outcomes are winsorized at the 1st and 99th percentiles. Salaries, total costs, and charges are reported in millions of dollars. Robust standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3. Bed-Adjusted Local Boundary Estimates

Outcome	5 km	10 km	15 km	20 km	25 km
FTE employment	586.544 (509.768)	843.409** (359.147)	538.660* (292.362)	453.564* (244.223)	157.115 (236.802)
Annual discharges	-1,134.486 (1,219.092)	-590.899 (795.052)	-832.611 (644.319)	-852.807 (547.997)	-575.515 (496.977)
Salaries, $ millions	79.370 (56.824)	87.399** (40.846)	45.061 (33.937)	45.780 (28.476)	29.487 (25.297)
Total costs, $ millions	131.932 (105.294)	162.220** (73.328)	117.663** (57.829)	100.988** (47.029)	62.276 (42.067)
Charges, $ millions	721.096 (542.596)	502.335 (396.332)	168.519 (304.975)	128.279 (256.474)	236.855 (227.089)
Wage-cost share	-0.008 (0.013)	0.003 (0.009)	0.007 (0.007)	0.009 (0.006)	0.008 (0.006)

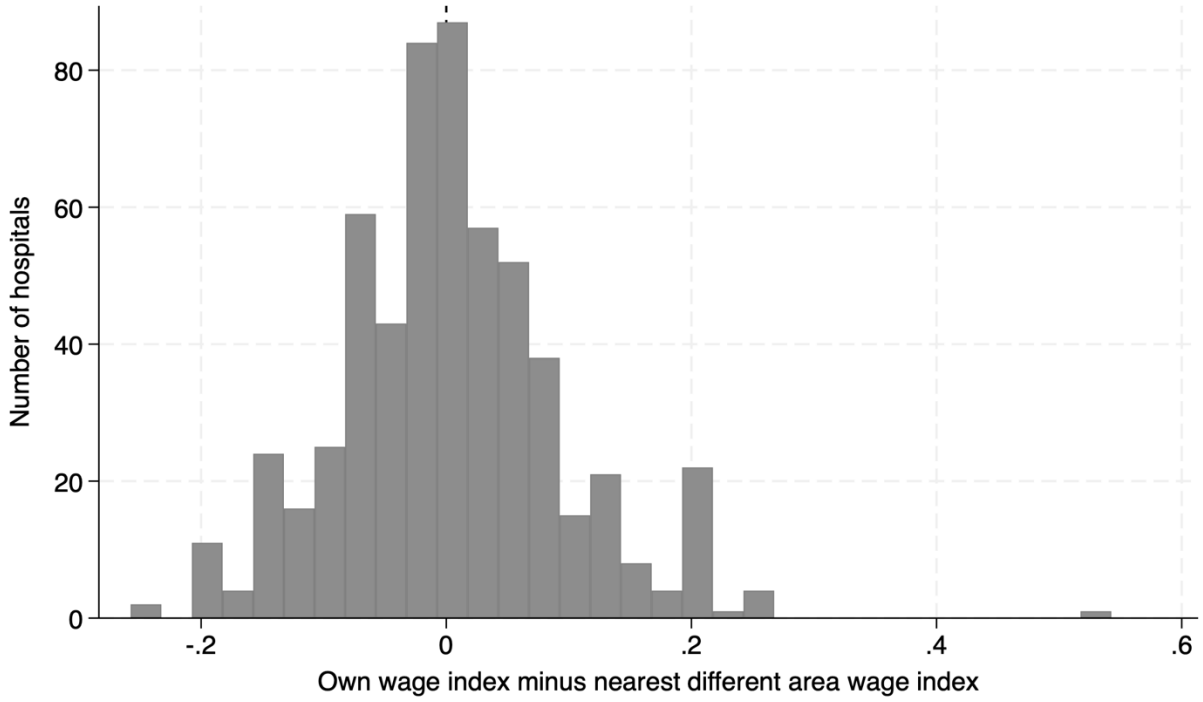
Notes. Each cell reports the coefficient on the higher wage-index side from a local linear boundary regression that controls for winsorized beds. Beds are excluded as an outcome row. The running variable is signed distance in kilometers to the nearest different wage-index area, with positive values indicating the higher wage-index side. Outcomes are winsorized at the 1st and 99th percentiles. Salaries, total costs, and charges are reported in millions of dollars. Robust standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 1. Wage Index Distribution



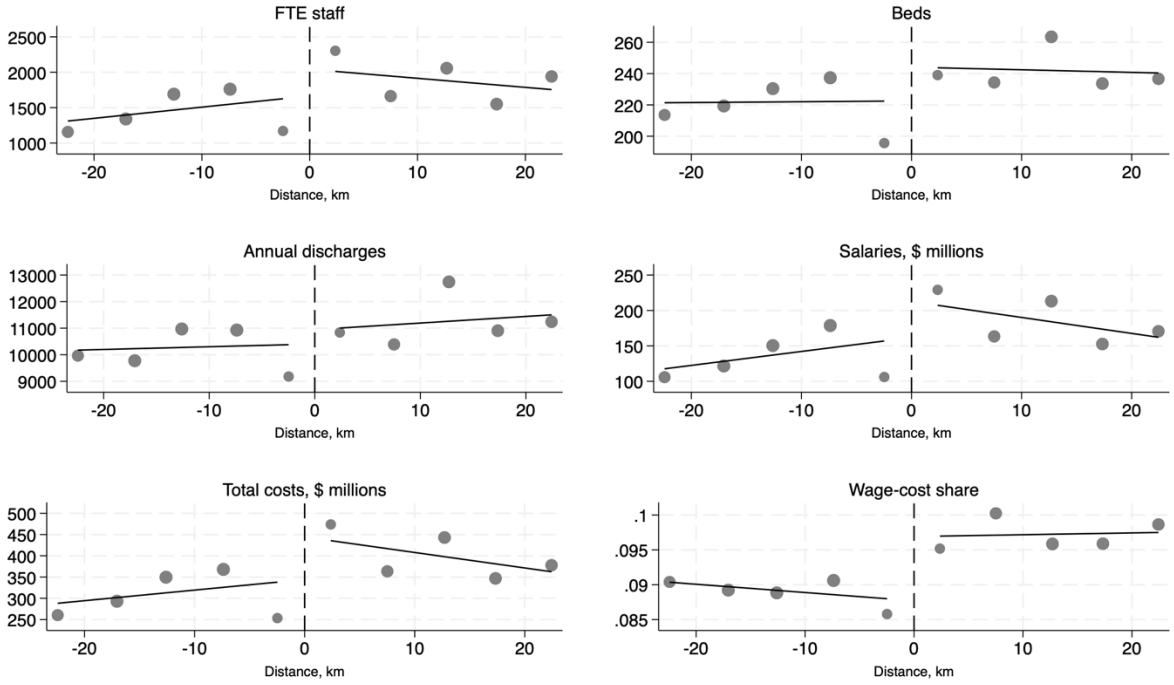
Notes. The figure plots the distribution of assigned FY 2026 hospital wage-index values in the matched analytic sample.

Figure 2. Boundary Wage-Index Gaps



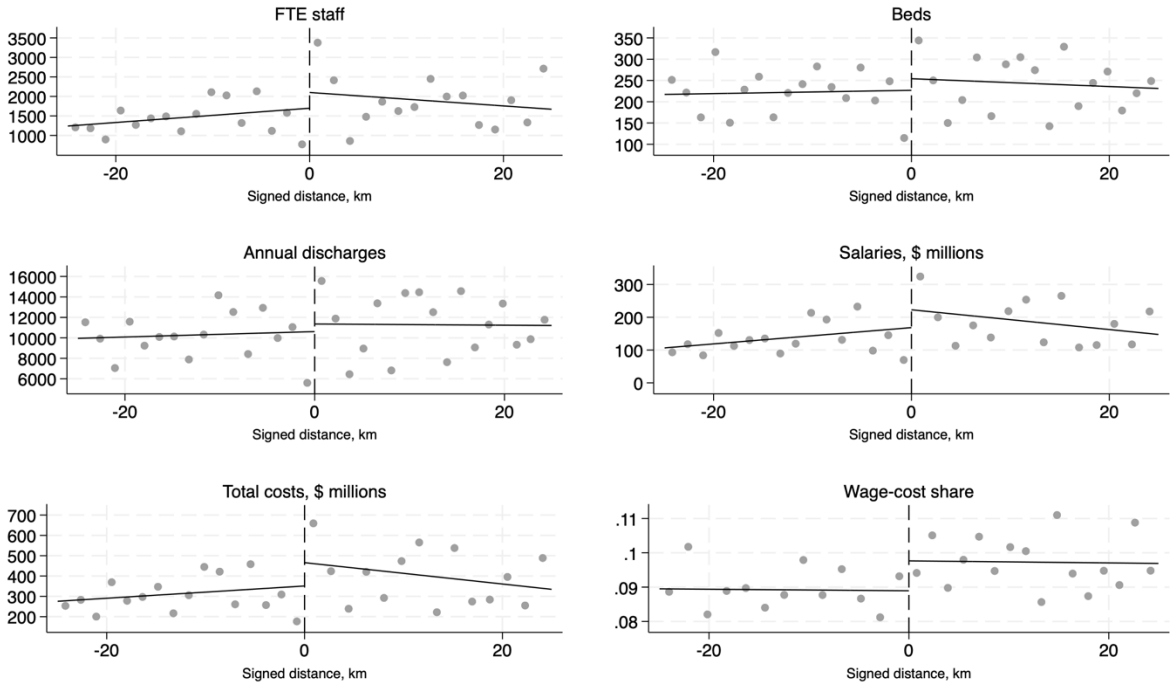
Notes. The figure plots the distribution of wage-index gaps within 25 kilometers of the nearest different wage-index area. The gap is the hospital's own wage index minus the wage index of the nearest comparison area.

Figure 3. Selected Outcomes Near Wage-Index Boundaries



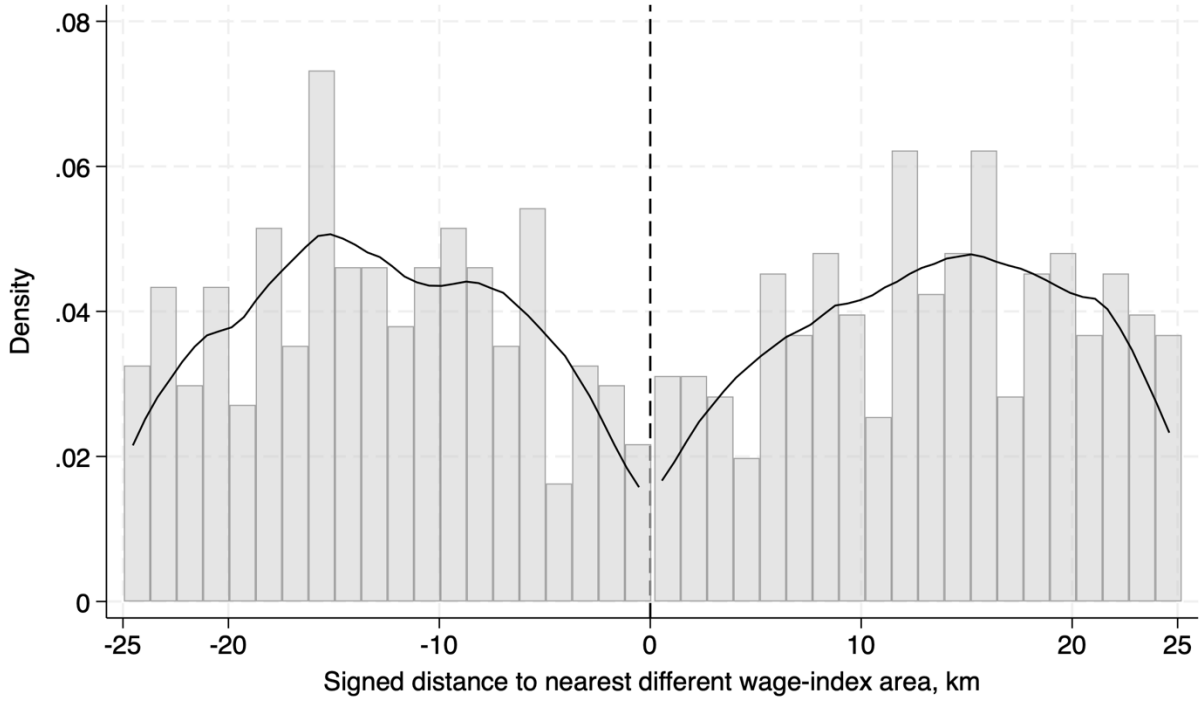
Notes. The figure plots binned means of six hospital outcomes against signed distance to the nearest different wage-index area. Positive distance indicates the higher wage-index side. Outcomes are winsorized at the 1st and 99th percentiles. Salaries and total costs are reported in millions of dollars. Points are weighted by the number of hospitals in each bin, and separate linear fits are shown on each side of the boundary.

Appendix A1. Data-driven RD Plots



Notes. The figure reports data-driven RD plots for six hospital outcomes using the same 25 kilometer local window as Figure 3. The running variable is signed distance to the nearest different wage-index area, with positive values indicating the higher wage-index side. The bins are selected using the optimal data-driven RD plotting procedure of Calonico et al. (2015). Outcomes are winsorized at the 1st and 99th percentiles. Salaries and total costs are reported in millions of dollars.

Appendix A2. Density of the Running Variable at the Boundary



Notes. The figure plots the density of signed distance to the nearest different wage-index area within 25 kilometers of the boundary. Positive distance indicates the higher wage-index side. The formal rddensity test produces a test statistic of -0.132 ($p = 0.895$).

Appendix A3. Ownership Balance Checks

Ownership category	5 km	10 km	15 km	20 km	25 km
Nonprofit	-6.151	9.509	7.529	1.040	-4.518
	(21.797)	(14.517)	(11.497)	(9.850)	(8.603)
	[75]	[200]	[328]	[451]	[563]
Proprietary	14.421	-7.541	2.160	7.329	4.782
	(15.793)	(11.906)	(9.581)	(8.328)	(7.133)
	[75]	[200]	[328]	[451]	[563]
Government	-8.270	-5.454	-7.689	-8.295	0.630
	(18.436)	(10.226)	(8.337)	(6.916)	(6.185)
	[75]	[200]	[328]	[451]	[563]

Notes. Each cell reports the coefficient on the higher wage-index side from the same local linear boundary regression used in the main analysis, but with ownership indicators as outcomes. Estimates and standard errors are reported in percentage points. The ownership categories are constructed from the CMS Hospital General Information file. Robust standard errors are in parentheses. Sample sizes are reported in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix A4. Same-State Boundary Estimates

Outcome	5 km	10 km	15 km	20 km	25 km
FTE employment	3,921.807**	1,843.685*	1129.712	809.407	325.918
	(1681.090)	(1058.457)	(786.070)	(634.109)	(550.182)
Beds	294.278**	80.930	34.450	21.154	10.822
	(132.986)	(84.046)	(63.310)	(51.626)	(44.536)
Annual discharges	12,909.485*	3275.723	1043.590	410.221	168.092
	(6592.036)	(4163.497)	(3097.905)	(2579.829)	(2237.207)
Salaries, $ millions	414.355**	198.826*	108.743	77.590	41.685
	(176.223)	(111.914)	(83.875)	(68.193)	(58.036)
Total costs, $ millions	788.835**	382.339*	248.409	180.356	101.876
	(344.946)	(214.210)	(158.881)	(127.508)	(108.788)
Charges, $ millions	3,802.584**	1475.526	618.675	361.199	373.943
	(1677.424)	(1003.717)	(755.339)	(610.588)	(519.959)
Wage-cost share	0.008	0.006	0.011	0.011	0.009
	(0.015)	(0.010)	(0.008)	(0.007)	(0.006)
N	60	162	274	386	486

Notes. Each cell reports the coefficient on the higher wage-index side from a local linear boundary regression at the indicated bandwidth, estimated on the subsample of hospitals whose nearest comparison county belongs to the same state. Comparisons in which the nearest different wage-index county is across a state line are excluded. The running variable is signed distance in kilometers to the nearest different wage-index area, with positive values indicating the higher wage-index side. Outcomes are winsorized at the 1st and 99th percentiles of the full analytic sample. Salaries, total costs, and charges are reported in millions of dollars. Robust standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix A5. Back-of-the-Envelope Magnitude Calculation

Outcome	Bandwidth	N	Mean Outcome	Mean abs. Gap	Boundary Estimate	Standard error	Implied Diff for a 0.01 Gap
FTE employment	5 km	76	1,737.163	0.076	3,673.990**	1,596.31	481.635
	10 km	202	1,723.685	0.066	1,756.148*	947.946	266.892
Beds	5 km	76	217.355	0.076	276.673**	124.373	36.27
	10 km	202	228.965	0.066	86.136	75.324	13.091
Annual discharges	5 km	76	10,009.724	0.076	12,300.279**	6,116.78	1,612.48
	10 km	202	10,428.055	0.066	3,556.63	3,719.30	540.521
Salaries, $ millions	5 km	76	167.612	0.076	391.511**	167.267	51.324
	10 km	202	170.168	0.066	183.935*	99.736	27.954
Total costs, $ millions	5 km	76	363.666	0.076	753.918**	326.003	98.834
	10 km	202	365.119	0.066	350.338*	192.646	53.243
Charges, $ millions	5 km	76	1,602.457	0.076	3,555.265**	1,568.62	466.071
	10 km	202	1,715.121	0.066	1,386.89	904.072	210.773

Notes. The table rescales the reduced-form boundary estimates by the mean absolute wage-index gap within each bandwidth. The last column reports the value of $\hat{\tau}_{S,b}\left(\frac{0.01}{\overline{|g|}_b}\right)$, where $\hat{\tau}_{S,b}$ is the Table 2 estimate for outcome S and bandwidth b. Only the 5 km and 10 km windows are displayed both to conserve space and to reflect the fact that the main reduced-form estimates are most informative in these local windows. *Salaries*, *total costs*, and *charges* are reported in millions of dollars. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.