WORKING PAPERS

PRICES FOR MEDICAL SERVICES VARY WITHIN HOSPITALS, BUT VARY MORE ACROSS THEM

MATTHEW PANHANS
TED ROSENBAUM
NATHAN WILSON

WORKING PAPER NO. 339

December 2018

PRICES FOR MEDICAL SERVICES VARY WITHIN HOSPITALS, BUT VARY MORE ACROSS THEM[1]

MATTHEW PANHANS (FTC), TED ROSENBAUM (FTC), AND NATHAN WILSON (FTC)

ABSTRACT

Using commercial claims for 2012-2013 from Colorado's All-Payer Claims Database, we examine how medical service prices vary for five hospital-based procedures and the complexity-adjusted inpatient price. We find that prices vary substantially in multiple dimensions. Our analysis indicates that there is significant price variation across payers for the same service in the same hospital. If prices converged to the lowest rate each hospital receives, commercial expenditures would fall by 10-20%. The share of overall price variation accounted for by hospitals variation tends to be even more substantial. For four out of six prices, we find that differences associated just with hospitals' metropolitan areas account for over 45% of the total variation. We observe substantial residual variation (17-50%) after accounting for factors specific to a given payer or provider.

Key words. Price dispersion, hospitals, payers, bargaining

INTRODUCTION

Prices for similar medical services vary substantially (Cooper, Craig, Gaynor, & Van Reenen, 2018). Economic theory suggests that these price differences could result from a range of factors including differences in costs, consumers' preferences, and market power (Grennan & Swanson, 2018). Research into the specific factors driving this price variation is important in developing policy, since it may suggest ways to lower overall healthcare expenditures without reducing the quality of care.

Using all-payer commercial claims data from Colorado, we document substantial price variation for six prices: five medical services and a complexity-adjusted inpatient price. We find variation exists across metropolitan areas, across hospitals within the same metropolitan area, and across payers within a hospital. Further, our analysis shows that lowering prices to the level paid to lowest priced facilities and by the lowest price payers would lead to substantial savings. For example, if all payers paid the price of the lowest price payer within every hospital, overall expenditures would decrease by 10-20% across the services we examine.

In order to understand the main factors leading to this variation, we decompose the variation in prices into differences across metropolitan areas, payers, and hospitals.[2] For four of the six prices that we consider, differences across metro areas account for over 45% of the overall variation in prices. Across hospital (but within metro area) variance was more substantial than across payer variance for four out of six prices we study and ranged from 8%-23% of the overall variation. Across payer variance ranged from 2%-50%, with four out of six prices below 11%. Interestingly, we still observe a substantial amount of variation across payer-provider pairs even after accounting for factors specific to a given payer or provider. For example, all of the

payers we observe sometimes pay more than a provider's median price and sometimes pay less than a provider's median price.

By examining the sources of price variation, we are able to generate intuition into where policymakers may most effectively devote attention. Our results about the large role played by geographies are consistent with the presence of cost or demand differences across areas, and fit with the large existing literature documenting such heterogeneity (Mays & Smith, 2009; Newhouse & Garber, 2013). The result that there are consistent price differences across providers could reflect differences in individual providers' costs (Schmitt, 2017) or market power (Cooper, Craig, Gaynor, & Van Reenen, 2018). Conversely, the existence of variation within providers is consistent with heterogeneity in demand and/or buyer power across payers (Trish & Herring, 2015; Ho & Lee, 2017). However, the large amount of variation that is not explained by consistent differences across geographies, payers, or providers suggests that negotiations between individual providers and payers contain other dimensions still to be effectively modeled.

New Contribution

Our results contribute to the growing literature in health economics on price dispersion for medical supplies (Grennan & Swanson, 2018; Grennan, 2013), pharmaceuticals (Starc & Swanson, 2018), and medical services (Cooper, Craig, Gaynor, & Van Reenen, 2018; Newman, Parente, Barrette, & Kennedy, 2016; Xu, et al., 2015; Dunn, Liebman, & Shapiro, 2017).[3] The previous papers looking at medical services have relied on data sets that do not include explicit information on the payer with which each patient is associated. In contrast, our all-payer claims data include payer identifiers exploitable in our analysis. This permits us to assess formally the share of price variation accounted for by both payer and provider components. Separately, our

research focuses on Colorado, which contains multiple metro areas and is not associated with dominant providers in the same way as other emerging research using all payer claims data (Craig, Ericson, & Starc, 2018).

STUDY DATA AND METHODS

The data for this study are medical claims for 2012-2013 from Colorado's All-Payer Claims Database (APCD). We use commercial claims from the individual and group markets.[4] Each claim includes information on the medical diagnosis, procedures performed, and the total allowed amount paid to the hospital. These expenditures reflect payments to the hospital; they do not include payments for the professional component of the services provided.

For each payer, we compute hospital-specific reference prices for five common and relatively homogeneous services (knee replacements, hip replacements, vaginal births, Caesarean section births, and MRIs) and for a complexity-adjusted measure of the average inpatient price.[5] To construct the reference prices, we restrict our sample to focus on services performed in general acute care hospitals, and drop the top and bottom 1% of prices to eliminate clerical billing errors or highly unusual medical events.[6]

For each service, we average the allowed amounts of claims associated with each hospital/payer pair to obtain a hospital/payer price. We only include the pair's price in our analysis if it was based off at least 50 admissions for the complexity-adjusted inpatient price measure, and at least 10 admissions for each of the procedures. Finally, for each service, we restrict our sample to include only hospitals and payers that are each part of at least two pairs. This restriction ensures that we can separate the contributions of each hospital and payer to price

variation.[7] Summary statistics for the data (after the outliers have been dropped) are in the first three columns of Exhibit 1.[8]

Using the sample of pair prices, we examine the distributions of prices paid by different payers to different hospitals. Specifically, we quantify the share of the overall variance in reference prices that comes from differences across metro areas, hospitals within metro areas, payers (e.g., high and low price payers), sorting (e.g., high price hospitals contract with high price payers), and residual variation unexplained by any of the above. We do this both using descriptive graphs and a formal variance decomposition, the details of which are described in Appendix B.

STUDY RESULTS

Price Variation across Services

We find that the prices paid for our reference services varied widely across payers and hospitals. This can be seen in Exhibit 1, which shows the weighted (by number of events) and unweighted average pair price in our data as well as its standard deviation. For all of the price series, the ratio of the standard deviation to the mean is at least 0.21. In other words, even for seemingly homogeneous services such as MRIs or uncomplicated vaginal births, payers are reimbursing different hospitals within the same state very different amounts.[9] Interestingly, this ratio is highest (0.41) for MRIs, likely the most homogeneous of the services we study.[10]

Price Variation within and across Hospitals

We find that hospitals are reimbursed at different rates for identically coded services. In other words, the price received by a given hospital may vary substantially for the same knee replacement, vaginal birth, etc. We demonstrate the magnitude of variation in Exhibit 2, which

shows the average price and the range of prices that hospitals in our sample receive for the five individual services we study.

Across all five of the services we study, the Exhibit shows that the range is frequently quite large within hospitals. However, the Exhibit also shows that there are wide differences in prices across hospitals, both between metro areas and within the same metro areas. These results imply that marginal cost or demand differences at the metro area level cannot explain all of the variation in prices.

Price Variation across Payers

In Exhibit 3, we show the frequency that each payer's price was above the median received by individual hospitals for each service. While some payers often have relatively lower prices than their rivals, there was no payer that always reimbursed above or below the median price across all services.

If prices simply reflect payers' relative bargaining leverage, and payers' bargaining leverage is constant across Colorado, then one would expect some payers to always have lower prices. The fact that we do not observe such consistency suggests that payer leverage varies across geographies, that payers do not apply their leverage consistently across all services, and/or that there are other factors that may lead to variation in prices.

Variance Decomposition

In our variance decomposition (Exhibit 4), we formally quantify the shares of overall price variation attributable to differences across payers, metro areas, hospitals (after accounting for metro area), and idiosyncratic differences across pairs. The shares associated with these different sources must sum to 100%, but can include a term for the hospital/payer covariance due

to the possibility for positive or negative sorting (e.g., high-price hospitals contracting more frequently with low-price payers).

We find substantial price variation across metro areas. For four out of the six services, the across metro area variation accounts for over 45% of the total variation in prices. After accounting for variation across metro areas, both payer and within metro area hospital variance are important, with hospital variance more substantial than payer variance for four out of six services. For four out of the five individual services, the share of the variation explained by differences across payers is under 11% and for the complexity-adjusted inpatient price, the share is 22%.

Finally, the portion of the variance attributed to hospital/payer pair specific factors is generally above 20%. The hospital/payer covariance term generally did not play a significant role in explaining the variation. Hip Replacements are an exception, but have the smallest sample size of any of the services we study.

Quantifying the Implications of Different Reductions in Variance

In Exhibit 5, we quantify the magnitude of different sources of price variation by comparing the current average prices of six different services to the average prices that would exist in four counterfactual scenarios of reduced price dispersion. In the first counterfactual, we compute the average prices that would be paid if all patients were kept in the same hospital, but shifted coverage to the lowest priced payer. In the second, we hold fixed each patient's payer, but calculate the average price that would be paid if the patients received care at the lowest priced hospital within their home metro areas.[11] In the third, we hold patients' payer constant, but shift them to their network's lowest priced hospital statewide. In the fourth, we calculate average prices using the lowest price observed in the data across both payers and providers.

Across all of the services we consider, we find at least a 10% reduction in average prices if all payers contracted at the price of the lowest priced payer within the same hospital. This is similar to the magnitude of price reduction if all patients were shifted to their payer's lowest priced hospital in a metro area. However, for five out of six prices, these reductions are substantially smaller than the reduction that would be obtained from shifting patients to each plan's lowest price hospital within the state (at least a 20% reduction).

Robustness

We conduct two robustness checks to see whether our results stem from unobserved sources of variation. Since we pool observations over a two-year period, one might be concerned that our results are driven by changes to contracts that occur during this two-year time frame. As we discuss in Appendix C, our results are qualitatively robust to using one year of data. Alternatively, one might worry that some of the dispersion in prices reflects different patient mixes in the hospitals. In Appendix D, we show that our results are robust to using individual claims data and risk adjusting by patient and treatment characteristics.

DISCUSSION

Many studies have found evidence that health care prices vary widely. Much of the prior focus has been on differences *across* hospitals. This paper represents one of the first to demonstrate that another significant driver of overall price dispersion is variation *within* hospitals.

When we compare the importance of the *within* versus *across* hospital variation, our descriptive results and formal variance decomposition show that a large share of the variation in prices is attributable to cross-metro area variation. For most of the services we study, eliminating

8

dispersion across hospitals throughout the entire state would produce double the savings compared to eliminating the dispersion across payers or across hospitals within a metro area.

When we focus on the *within* hospital variation, we find meaningful variation in prices and some evidence suggesting that certain payers tend to pay higher/lower rates. This would support the views of some papers in the economic literature suggesting that stronger payers possess greater bargaining leverage (Hemphill & Rose, 2018; Ho & Lee, 2017). However, since we find that there is no universally "high-price" or "low-price" payer, our results suggest caution in viewing payer size as the major driver of price differences in health care. Rather, the variation in negotiated prices across payer-hospital pairs suggests that many factors affect the outcome of these interactions. Explaining this variation is a fruitful area for further research.

CONCLUSION

The prices of seemingly similar health care services vary widely, even within a hospital. Our analysis illustrates substantial price variation due to both sides of the table in payer-hospital price negotiations. In addition, our results reinforce the importance of determining the reason for variation in prices across metro areas, providers, and payers. For example, if prices paid by each payer converged to the lowest rate each hospital receives, expenditures would fall by 10-20%.

REFERENCES

Baye, M. R., Morgan, J., & Scholten, P. (2004). Price dispersion in the small and in the large: Evidence from an internet price comparison site. *The Journal of Industrial Economics, 52*(4), 463-496.

Cooper, Z., Craig, S., Gaynor, M., & Van Reenen, J. (2018). *The price ain't right? Hospital prices and health spending on the privately insured.* Cambridge: NBER.

Craig, S. V., Ericson, K. M., & Starc, A. (2018). *How Important Is Price Variation Between Health Insurers?* National Bureau of Economic Research.

Dunn, A., Liebman, E. B., & Shapiro, A. (2017). *Decomposing medical-care expenditure growth.* National Bureau of Economic Research.

Farrell, J., Balan, D. J., Brand, K., & Wendling, B. W. (2011). Economics at the FTC: Hospital Mergers, Authorized Generic Drugs, and Consumer Credit Markets. *Review of Industrial Organization, 39*(4), 271-296.

Gaynor, M., Ho, K., & Town, R. (2015). The Industrial Organization of Health Care Markets. *Journal of Economic Literature, 53*(2), 235-284.

Grennan, M. (2013). Price discrimination and bargaining: Empirical evidence from medical devices. *American Economic Review, 103*(1), 145-177.

Grennan, M., & Swanson, A. (2018). Diagnosing Price Dispersion. *Wharton working paper*.

Hemphill, C. S., & Rose, N. (2018). Mergers That Harm Sellers. *Yale Law Journal, 127* .

Ho, K., & Lee, R. (2017). Insurer Competition in Health Care Markets. *Econometrica, 85*(2), 379-417.

Kaplan, G., & Menzio, G. (2015). The morphology of price dispersion. *International Economic Review, 56*(4), 1165-1206.

Mays, G., & Smith, S. (2009). Geographic variation in public health spending: correlates and consequences. *Health Services Research, 44*(5), 1796-1817.

Newhouse, J. P., & Garber, A. M. (2013). Geographic variation in health care spending in the United States: Insights from an Institute of Medicine report. *JAMA, 310*(12), 1227-1228.

Newman, D., Parente, S., Barrette, E., & Kennedy, K. (2016). Prices for common medical services vary substantially among the commercially insured. *Health Affairs, 35*(5), 923-927.

Schmitt, M. (2017). Do hospital mergers reduce costs? *Journal of Health Economics, 52*, 74-94.

Starc, A., & Swanson, A. (2018). What do private firms bring to the (bargaining) table? Vertical structure and negotiated prices of pharmaceuticals. *working paper*.

Trish, E. E., & Herring, B. J. (2015). How do Health Insurer Market Concentration and Bargaining Power with Hospitals Affect Health Insurance Premiums? *Journal of Health Economics, 42*, 104-114.

Xu, X., Gariepy, A., Lundsberg, L., Sangini, S., Pettker, C., Krumholz, H., & Illuzzi, J. (2015). Wide variation found in hospital facility costs for maternity stays involving low-risk childbirth. *Health Affairs, 34* (5), 1212-1219.

EXHIBITS

	Number of Providers	Number of Payers	Number of Pairs	Number of Events	Average Price (Weighted)	Average Price (Unweighted)	Standard Deviation of Prices
C–Section	20	7	61	2,881	$10,277	$11,309	$2,655
Hip Replacement	8	4	20	534	$27,554	$28,165	$6,015
Inpatient Price	32	8	150	79,292	$12,047	$13,878	$4,907
Knee Replacement	10	5	25	1,301	$26,688	$32,628	$9,824
MRI	17	6	47	1,394	$1,129	$1,162	$477
Vaginal Birth	28	7	103	8,372	$5,691	$6,065	$1,322

Exhibit 1. Summary statistics of prices for services in study sample. Source/Notes: SOURCE Authors' analysis of claims data from the Colorado All Payer Claim Database.

Exhibit 2. Range of payer prices at providers. Source/Notes: SOURCE Authors' analysis of claims data from the Colorado All Payer Claim Database. NOTES CBSAs and providers are anonymized. The point in the middle is the weighted average price at that provider. The colors represent the CBSA of the provider.

Exhibit 3. Fraction of providers at which each payer is above median price. Source/Notes: SOURCE Authors' analysis of claims data from the Colorado All Payer Claim Database. NOTES Payers are presented in the order of their size, with A the largest payer.

Exhibit 4. Decomposition of price variation by service or complexity-adjusted inpatient price. Source/Notes: SOURCE Authors' analysis of claims data from the Colorado All Payer Claim Database.

Exhibit 5. Changes in average prices paid under counterfactual scenarios. Source/Notes: SOURCE Authors' analysis of claims data from the Colorado All Payer Claim Database. NOTES Numbers are rounded to the nearest hundred.

NOTES

---

[1] The data used in this study are from the Center for Improving Value in Health Care (CIVHC), Administrator of the Colorado All Payer Claims Database. Information about accessing the data is available at www.civhc.org.

[2] By "prices", we mean the claim allowed amounts actually paid to providers for services, including both the patient and payer contributions. We use the term "payer" to refer to the entity that constructs and manages the provider network associated with an individual patient's commercial insurance plan. We do not have data to determine whether a plan is self-insured.

[3] There is an abundant literature in economics on price dispersion in other settings (Kaplan & Menzio, 2015; Baye, Morgan, & Scholten, 2004).

[4] We exclude claims from, e.g., the Medicare and Medicaid populations.

[5] See, e.g., Cooper et al. (2018). To compute the complexity-adjusted inpatient price measure, we divide each price by the admission's associated DRG average resource weight. We use the MSDRG weights available from CMS that reflect average resource used to treat Medicare patients in each DRG. Appendix A provides more details on the specific MSDRG codes used.

[6] This follows Cooper et al. (2018).

[7] Appendix A provides more details on our processing of the claims data.

[8] Appendix C discusses the distribution of prices within each provider-payer pair.

[9] We note that there is also significant variation within a payer/hospital pair at the individual level. This finding corroborates Cooper et al. (2018). These ratios are computed using the unweighted mean.

[10] The magnitude of variation is larger than, but similar to, that documented for medical supplies in recent research, which found coefficients of variation ranging from 9% to 35% (Grennan & Swanson, 2018).

[11] We use core-based statistical areas (CBSAs) to proxy for local metro areas. Details on CBSAs are available at: https://www.census.gov/geo/reference/gtc/gtc_cbsa.html. We are not

asserting that these areas represent relevant antitrust markets as described in the *Horizontal Merger Guidelines* issued jointly by the Federal Trade Commission and Department of Justice. There are 11 CBSAs in our dataset.

A. Data Appendix

We restrict the sample to services paid for by commercial insurance plans in 2012-2013. A payer is defined as the parent insurance company, encompassing all of that firm's commercial insurance products. We also restrict the sample of providers to include only general acute care hospitals identified in the American Hospital Association directory. Furthermore, we kept only claims that listed a hospital's primary NPI, thus excluding claims that were generated by hospital emergency departments, home health units, and behavioral health programs. This was verified by checking the NPI description as listed on the NPPES NPI Registry. In order to account for price variation due to clerical billing errors or unusual medical events, we drop the top and bottom 1% allowed amount claims by service, as is also done in Cooper et al.

For the inpatient sample, we create a treatment intensity weighted inpatient price by dividing each price by the DRG average resource weight. Thus, only inpatient claims with an assigned MS-DRG code were included in the inpatient sample. We further drop any inpatient admissions with a DRG that has a frequency of less than 1%.

For each hospital-payer pair, we only include the pair's price in our sample if it had at least 50 admissions, for the inpatient price, or at least 10 procedures. Further, we only include hospitals with price pairs for least two payers and payers with price pairs for at least two hospitals. We do this sequentially, starting with all payer and hospital pairs, dropping hospitals and payers by the criteria above, and iterating until convergence. We do this on a service-by-service basis (treating our complexity-adjusted inpatient price as a service).

We identify straightforward, non-complicated events for our samples of procedures based on the codes given in the table below. For MRIs, we selected claims only for those of imaging of the lower extremities, based on CPT code 73721. Additionally, we kept MRI claims only if the scan was the only line on the claim.

Procedure	Diagnosis Code	and	Procedure Code	CPT-4
Hip Replacement	APR-DRG 301 or MS-DRG 470		ICD9 81.51 or 81.52 or ICD10 0SR903Z, 0SR904Z, 0SRB04Z	
Knee Replacement	APR-DRG 302 or MS-DRG 470		ICD9 81.54 or ICD10 0SRC0J9 or 0SRD0J9	
Vaginal Birth	MS-DRG 775		ICD9 73.59 or 75.69 or ICD10 10E0XZZ	
Cesarean Section	MS-DRG 766		ICD9 741 or ICD10 10D00Z1	
MRI				73721

B. Derivation of Variance Decomposition

Suppose a linear model of the form:

$$y_{ij} = \alpha_i + \gamma_j + \varepsilon_{ij}$$

where in our context, y_{ij} would denote the total allowed amount paid to hospital i and insurer j, α_i are hospital fixed effects, γ_j denotes insurer fixed effects, and ε_{ij} denotes the error term. After estimating the parameters of the model by least squares, the covariation with the dependent variable can be decomposed as follows by the variance property for linear combination and the fact that the error term is uncorrelated with the regressors:

$$var(y) = var(\alpha + \gamma + \varepsilon) = var(\alpha) + var(\gamma) + var(\varepsilon) + 2cov(\alpha, \gamma)$$

We then further decompose the hospital variation into geographic (here, CBSA) variation and net-of-geography hospital variation by regressing the hospital fixed effects on CBSA fixed effects, which allows us to allocate hospital variation into those two components. Specifically, we regress:

$$\alpha_i = \delta_{g(i)} +$$

where $\delta_{g(i)}$ is a fixed effect for hospital i's CBSA. Thus analogously to before, we can derive that:

$$var(\alpha) = var(\delta) + var(\eta)$$

where the first term indicates the share of the hospital fixed effect variation attributable to geography, while the second term is the share attributable to hospitals net-of-geography.

Therefore, the complete variance decomposition is:

$$var(y) = var(\delta) + var(\eta) + var(\gamma) + var(\varepsilon) + 2cov(\alpha, \gamma)$$

Using this expression, we obtain the variance decomposition in Exhibit 5, where the components are shown as percentages of the total variance in price ($var(y)$).

C. Robustness of Single Year Analysis

In our main data sample, the share of events with a price equal to the modal price within the hospital-payer-procedure is 64% for MRIs, 30% for Cesarean sections, 32% for vaginal births, 32% for knee replacements, and 38% for hip replacements, suggesting that a fairly homogenous set of procedures have been identified. Part of the reason why a higher share of prices do not equal the modal price is because our main sample uses events from 2012-2013, and presumably many pricing contracts will change over that time period. When we restrict the data to only one year, we find as expected that the share of events with prices equal to the modal price is even higher. Specifically, the share of events with a price equal to the modal price within a hospital-payer-procedure year is 80% for MRIs, 45% for C-sections, 49% for births, 51% for knee replacements, and 63% for hip replacements.

This section thus reproduces the main results using only events occurring in calendar year 2013. We find largely the same patterns. Note however that we drop hip replacements from the decomposition analysis, because using only one year of data leaves us with too few hip replacement events to find meaningful results.
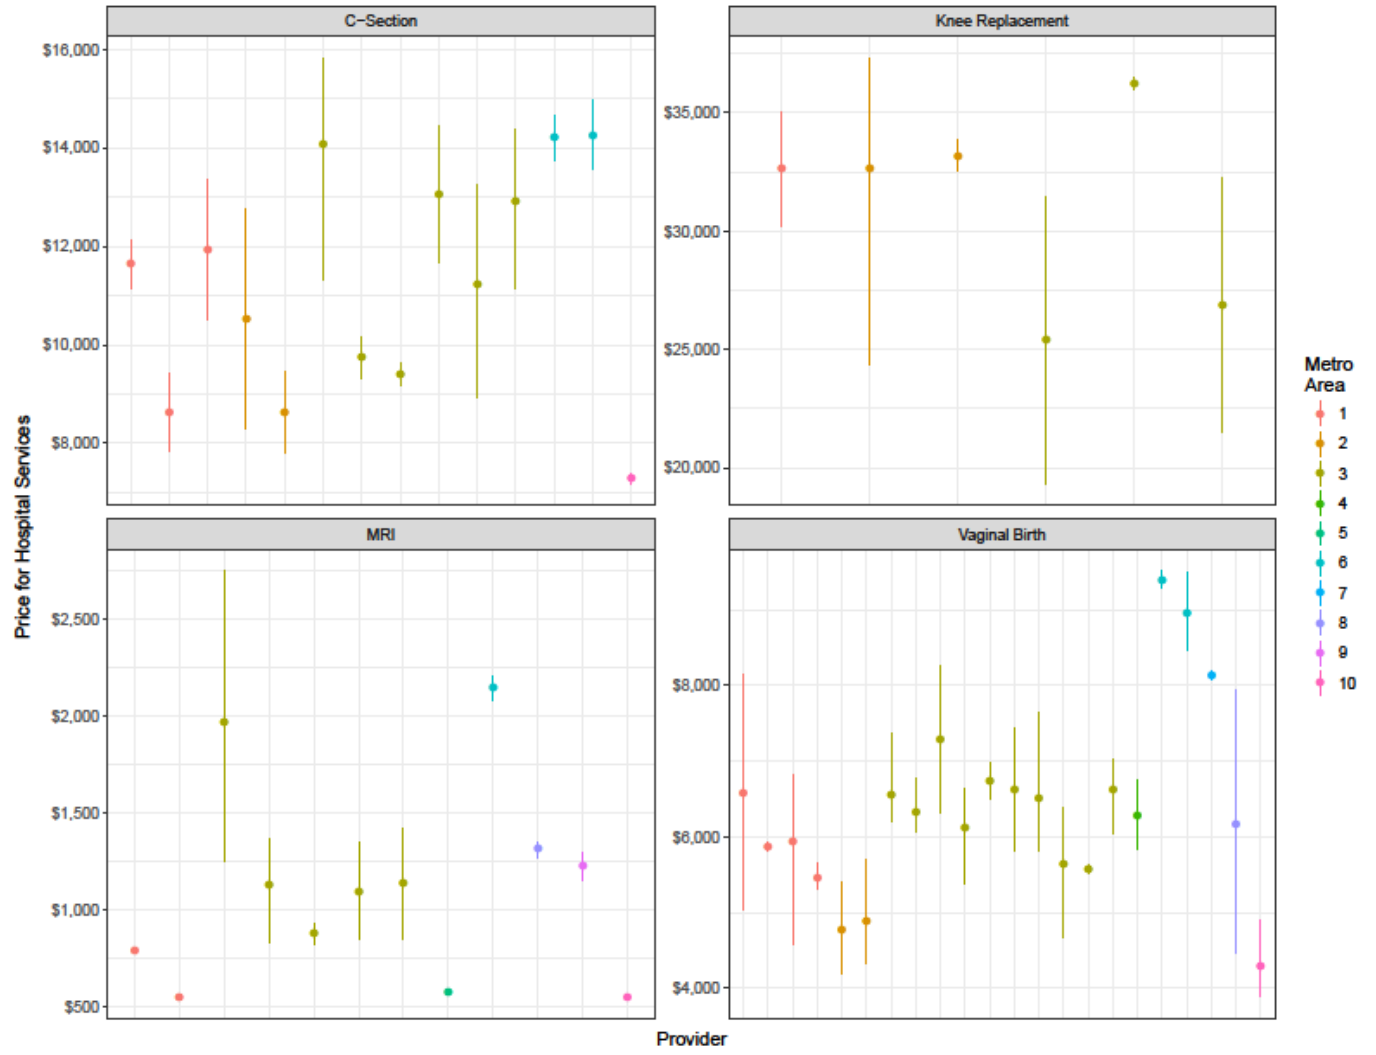
Exhibit A1. Robustness of Exhibit 2 to only using 2013

Exhibit A2. Robustness of Exhibit 3 to only using 2013

Exhibit A3. Robustness of Exhibit 4 to only using 2013

Exhibit A4. Robustness of Exhibit 5 to using only 2013

D. Robustness of Risk-Adjustment on Individual Characteristics

As a robustness exercise, we create price indices that are risk-adjusted using individual person characteristics (gender and age) and claim details (length of inpatient stay and number of claim lines on claim).

Specifically, in the case of the inpatient sample, we run an event-level regression:

$$p_{i,j,h,d} = \alpha_h + \gamma_j + X_{i,d}\beta + \phi_d + \varepsilon_{i,j,h,d}$$

where $p_{i,j,h,d}$ denotes the payment made for event i from insurer j to hospital h for DRG d, α_h are hospital fixed effects, γ_j are insurer fixed effects, $X_{i,d}$ are event-level characteristics (age, gender, length of stay, number of lines on claim), ϕ_d are DRG fixed effects, and $\varepsilon_{i,j,h,d}$ is the stochastic error term.

After estimating the regression parameters by OLS, we predict the hospital-insurer prices $\hat{p}_{j,h}$ at the average patient characteristics $\bar{X}$ and the sample mean basket of DRG codes $\bar{\phi}$ across the entire sample:

$$\hat{p}_{j,h} = \hat{\alpha}_h + \hat{\gamma}_j + \bar{X}\hat{\beta} + \bar{\phi}$$

This yields a measure of a hospital-insurer pair risk-adjusted price. Note that variation in this adjusted price across hospital-insurer pairs comes only from the insurer and hospital fixed effects. We then replicate all of the graphs using this data set, and find similar effects. The results from this follow.
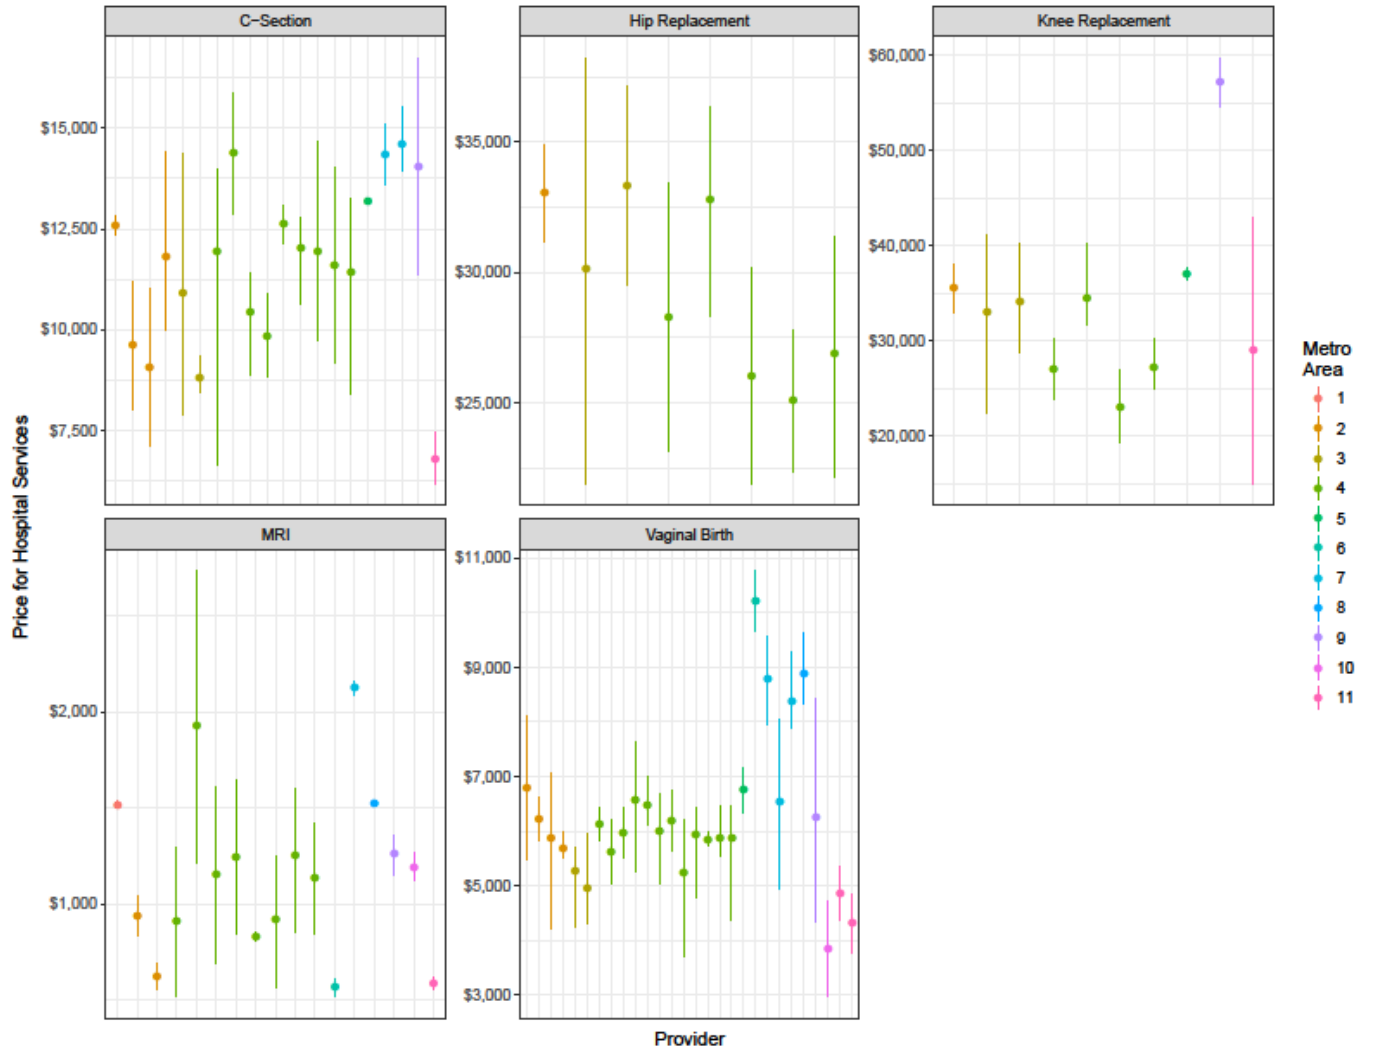
Exhibit A5. Robustness of Exhibit 2 to using risk-adjusted prices

Exhibit A6. Robustness of Exhibit 3 to using risk-adjusted prices

Exhibit A7. Robustness of Exhibit 4 to using risk-adjusted prices

Exhibit A5. Robustness of Exhibit 5 to using risk-adjusted prices