Design and Application of HSRV Log Weights

GRP-395
Publication date March 29, 2023

Table of Contents

Design and application of HSRV log weights

The purpose of relative weights is to provide a measure of relative resource use. For example, the charge-based method employed by CMS (Centers for Medicare & Medicaid Services) standardizes charges using payment formulae adjustments specific to the CMS IPPS (inpatient prospective payment system). This standardization aligns the relative charges across the CMS concept of efficient cost. The log Hospital-Specific Relative Value (HSRV) method takes advantage of the mathematical properties of logarithms within a multivariate linear regression model to achieve an analogous outcome by computing standardizing charge adjustment factors for hospitals that minimize the squared deviation of the interaction of charge and relative weight across hospitals within a Diagnosis-Related Group (DRG) system. While analogous to other standardization models, the log HSRV method eliminates bias within the relative weights associated with high-cost hospitals, thereby forcing supplemental payment adjustments in support of hospital missions to be transparent.

Relative weights: Background and objectives

Prospective payment systems (PPS) are designed to provide a pre-determined fixed payment for a unit of service. For an IPPS, the unit of service has evolved to become the inpatient discharge. Classification software for DRGs take information from patient claims data to pool hospital discharges within discrete, mutually exclusive categories. DRGs provide a standardized unit of service for hospital inpatient activity.

Standardizing units of service is only the first step in computing a fixed payment amount. Final payment requires calculating a price for the unit of service. The usual requirement for IPPS designs is that payment be based upon efficient cost. There are two related components in determining what constitutes efficient cost. The first, and most obvious, is the direct cost of providing care to patients. If all other factors are equal, then the final payment for a discharge would be based upon the average cost for a DRG. This would simplify calculation by merely requiring an index of relative values (relative weights) descriptive of DRG cost differentials and a base rate (standard payment amount) that projects total expenditure to meet payer budget. The system would allocate available resources by weighted DRG, so matching the share of cost with the share of payment. All factors are not equal, however.

Hospitals have different cost structures. Some variation is predictable and may result from regional differences (wage index, urbanicity); mission-related activities (teaching, providing care to the uninsured/under-insured); spearheading new technologies and research; providing regional specialty on-call services (trauma, burn centers), to give some examples. Many of these items have been explicitly recognized within payment systems as warranting separate additional payment. Typically these supplemental payments are made transparently and capped at certain empirical levels such that only the portion of cost deemed reasonable or permissible is passed through as payment. Another (unwanted) source of cost variation is relative inefficiency, and this is not suitable for payment adjustment. The concept of paying for efficient cost is therefore achieved by paying a patient-level rate to cover direct variation in the efficient cost per unit of service and hospital-level adjustments that reflect the unique characteristics of hospitals as allowed by policy.

In summary, relative weights are required to describe patient-level differences in unit cost across DRGs.

Weighting bias

The practical determination of relative cost requires moving beyond theory. While hospitals may report aggregate costs, this is not connected to an individual patient in a way that would enable the average cost per unit of service to be readily calculated. What is submitted to payers for each patient is a summary of the charges for care received in the format of routine standardized claims data. The use of claims-based charges is therefore the most practical solution to identifying routine patient-level cost. To achieve the objective of calculating relative differences in the cost of care for individual units of service, we therefore have to convert charges to underlying cost or be convinced that the direct use of charge data acts as a sufficiently close proxy to the underlying cost.

In practice we face a number of structural problems with charges in claims data that require attention. Cost-to-charge ratios (CCRs) are not uniform across hospitals. CCRs are not uniform within hospitals. Hospitals with high-cost structures tend to provide disproportionate amounts of some service units (DRGs). Hospitals categorically excluded from the Medicare PPS (children's hospitals) tend to have a higher level of x-inefficiency within their cost structure and dominate specific service units (DRGs).

The net effect of these structural claims data issues is that a measurement of relative cost that uses charges as a proxy without correction will tend to overstate the relative costliness of high markup service lines and those DRGs that have concentration in high-cost facilities. The resulting payment bias has been shown to result in improper incentives to provide services with higher markups (typically newer services[1] due to the effect of "sticky prices" in chargemaster maintenance), greater profitability in service lines favored by high-cost (typically academic medical center or AMC) hospitals, and an inadvertent redirecting of additional money through this reweighting to high-cost hospitals in support of their mission or inefficiency.

In summary, charge markups and the uneven distribution of case types lead to bias in relative weights that are based upon average charges.

Some history

Corrections based upon using relative cost were initially opposed in part because of technical complexity and in part due to the perceived issue of compression.[2] Compression is another payment bias whereby higher-cost cases are paid less than their fair ratio of cost, and lower-cost cases more. Historically the three main sources of compression have been

- Differential markups: High-cost services are marked up less due to sticker shock.
- Data incompleteness: A lack of coding completeness groups lower-cost cases together with higher-cost cases causing the mean weight of two different populations to converge.
- Classification limitations: Essentially the same effect as data incompleteness but requiring the DRG classification system to be modified.

As coding and classification systems have evolved from the early 1980s IPPS implementation, the fears of compression have drifted predominantly towards differential markups. In 2006, as CMS attempted to move the relative weighting method towards cost-based relative values in response to Medicare Payment Advisory Commission's (MedPAC's) report showing harmful structural changes in the healthcare delivery system were in part the result of using charge-based weights,[3] the specter of compression was a central plank in the concerted effort to stymie reform. In fact, industry reaction to the scope of proposed changes caused CMS to pull back from full-scale implementation of department-level HSRVs along with matching reform of Indirect Medical Education (IME) and Disproportionate Share Hospital (DSH) coefficients. The resulting change for FY2007 was a comparatively minor reform to the existing weighting method that extended the existing charge standardization using payment formulae variables to include a handful of national department-level cost-to-charge ratios.[4]

Since that time, attempts to address "charge compression" in relative weights have been led by academic medical centers as well as device and pharmaceutical manufacturers that would benefit from increasing relative weights for services that they are most closely associated with. These have been somewhat successful with a number of additional CCR divisions being introduced. The weighting method should not, however, be understood in isolation. The change to MS-DRGs moved money in one direction (to higher-severity hospitals like the AMCs); the use of CCRs moved money back the other way. Adding CCRs will redirect money back to AMCs. Both DRG and weighting reforms impact what can be empirically measured as justifiable financial support for IME and DSH since they are usually specified as residual factor regressions.

Having successfully introduced Medicare severity-adjusted DRGs (MS-DRGs), reclaiming monies for code creep, and transitioned to a version of cost-based weights, dealing with the PPACA and DSH reform for Medicare and Medicaid, all while championing the adoption of ICD-10, CMS and MedPAC understandably have not returned to the issue of bias in weight methods. This does not, however, mean that the CMS approach settled upon after 2007 is a preferred solution. There are significant limitations for generating weights in this way:

- They require standardization variables applicable to IPPS hospitals. Medicaid programs have many PPS-exempt hospitals for cancer, children critical access, etc.
- They utilize CCRs to impute operating and capital charge shares for adjustment while simultaneously acknowledging limitations in their use.
- They do not adjust for cell dominance by high-cost hospitals, despite evidence to the contrary.
- They provide no adjustment for inefficient provider cost levels.

The charge standardization method is flawed but is a sufficient political compromise for CMS to maintain the IPPS. It is not, however, suitable for extension beyond the CMS IPPS. In reality the cost-based HSRV approach touted by MedPAC was more theoretically pure, removing more of the bias in weights and requiring more transparency in funding high-cost hospitals (having to set empirical IME or DSH levels rather than bury profit in DRG weights), but operationally flawed. More precisely, while HSRV achieved greater standardization and was computationally feasible, the requirement to link cost report data at the cost center-level increased time lags and reduced accuracy due to the need to link detailed hospital cost reporting and charge generating practices. This resulted not only in delaying data availability by two years but also requiring acceptance of CCRs that ranged from 0.001 to 10. Despite the very public battle over HSRV cost centers, the disputes primarily concerned cost centers, not HSRV. Since bias and data complexity issues remain in weight setting, numerous agencies (for example, CMS for long-term care weights, Maryland's Health Services Cost Review Commission for all-payer rate setting, and New York Medicaid for inpatient rates) use the HSRV standardized charge method to remove bias from the resulting relative weights.

In summary, the use of HSRV in various forms has been widely debated and is an accepted technique for removing bias at the hospital level. Claims of compression have emanated from stakeholders fearing reduced payments principally due to viewing weights in isolation of other payment adjustments that support reasonable cost in high-cost hospitals. Cost center adjustments have proven more contentious due to poor data quality and delays in availability. CMS methods have limitations and represent a holding pattern. CMS standardization cannot be extended to non-IPPS hospitals (children's hospitals, cancer hospitals, critical access hospitals, etc.).

Hospital-specific relative values (HSRVs)

HSRV was first postulated by Vertrees and Pettengill as a solution to the issue of hospital-induced bias in relative weights prior to the implementation of the CMS IPPS.[5] As a general description, HSRV seeks to solve a simple logic problem. If hospitals have different cost structures and markups, then the average of those reported values will bias the average imputed relative cost of DRGs across hospitals. If instead we standardize the relative costliness of hospitals, then when we average the imputed relative cost of DRGs, we should have avoided bias. The catch, however, is that to create a standardization factor for a hospital, we must account for the different DRG mix which first requires determining the relative weights to be used.

To get around the circular logic that requires knowing the relative weight of DRGs before computing the standardization factor to help calculate the relative weight, an iterative approach was proposed. This is summarized in six steps:

1. Calculate the mean charge for all cases classified in the same DRG.
2. Calculate the overall mean charge across all DRGs.
3. Calculate the relative weights by dividing the mean charge for each DRG by the overall mean charge.
4. Calculate the hospital mean charge and case mix index (CMI) based on the relative weights computed in step three.
5. Adjust the case charge amount for each patient by dividing by the hospital's mean charge and multiplying by the hospital's CMI.
6. Repeat steps 1–5 until the relative weights are stable using the adjusted charge amount from step 5.

Each iteration adjusts hospital charges in proportion to average hospital charge and case mix calculated across all DRGs in all hospitals. When the average of the hospital standardized charges (standardized to relative values) equals the average of the hospital's DRG weights for each hospital, the relative weights have converged upon a solution. The result is a set of relative weights that approximate differences between DRG cost if relative values were computed for each hospital independently without regard to the hospital's relative charge level. There are other formulations that achieve the same end point, but the method will converge on the same root weights regardless of starting point as long as the weights are iterated to convergence.

As with other weighting methods, calculations can be sensitive to data trimming. The HSRV method is also adaptable for use as a geometric algorithm by using the log of charges to provide a more stable measure of central tendency in the weights. Given that most charge distributions are lognormal, and that outlier policy protects the high-end claims (often an issue for double counting within weights that also troubled MedPAC), it makes sense to improve stability of relative weights by centering HSRV weights around the geometric mean. In practice the difference between weights that are derived from arithmetic and geometric means tends to be negligible with larger data sets with more significant differences occurring when weaker trims are employed and with low-volume DRG cells.

In summary, the HSRV method eliminates bias in the relative weights that results from differences in hospital cost structure or markups.

Log HSRV

The log HSRV method was developed in part to increase the accessibility of understanding around the iterative adjustments in the HSRV method. Transforming the relative charge amounts into unitless ratios often leaves stakeholders perplexed that something untoward is occurring. To counter this reaction, the log HSRV method leverages modern software packages that enable large multivariate regression by using a simple mathematical relationship to achieve an HSRV solution. While the method closely approximates the iterative HSRV solution, it also brings some additional beneficial properties that are discussed below after the method explanation.

The iterative HSRV method contains two key ingredients:

- All charges in a hospital are ultimately standardized by a single coefficient unique to that hospital.
- Relative weights are calculated such that the case-mix adjusted standardized charge per case (CPC) is equal for all hospitals.

To understand the log HSRV approach, first let C_i be the charge for a case i. Each case i will have assigned one, and only one, of n APR DRGs D and have a single hospital J, of k total hospitals, for the case. Hospital J can potentially be the same hospital differentiated by time period, allowing older claims data to be introduced without the need to adjust for inflationary effects (i.e., treats the hospital as a unique instance). The requirement is simply that a single hospital within a set time period be identified by a unique identifier.

The charge for case i (C_i) is considered to be the product of the APR DRG ($D_{n,i}$) within which it is classified and the hospital ($J_{k,i}$) reporting the admission. Since the n APR DRGs D are mutually exclusive, they may be considered a binary variable that is either present 1 or absent 0 but not both. The k hospitals J may be considered similarly. This reduces the problem of calculating the relative weight W for the n APR DRGs to finding the best fit for the pattern of case charges i that are distributed across n APR DRGs and k hospitals, such that W provides a weight for APR DRG n and F provides a consistent hospital specific factor for hospital k across all APR DRGs. The hospital-specific charge adjustment factor F for hospital j serves the same case-mix standardization function as the iterative HSRV approach.

The equation defining this problem may be written as:

$$C_i = \omega_n D_{n,i} * \varphi_k J_{k,i} + \varepsilon_i$$

Where:

$$D_n = \begin{Bmatrix} 1\ if\ the\ APR\ DRG\ of\ case\ i\ is\ APR\ DRG\ n \\ 0\ if\ the\ APR\ DRG\ of\ case\ i\ is\ not\ APR\ DRG\ n \end{Bmatrix}$$

$$J_k = \begin{Bmatrix} 1\ if\ the\ hospital\ for\ case\ i\ is\ hospital\ k \\ 0\ if\ the\ hospital\ for\ case\ i\ is\ not\ hospital\ k \end{Bmatrix}$$

The charge for case i is considered to be the product of a consistent APR DRG charge ω across all hospitals and a consistent hospital base charge factor ϕ across all APR DRGs with residual error ε for case i. The consistent APR DRG charge ω is straightforwardly converted to a more traditional relative weight W by selecting an appropriate index value.

In calculating relative weights from cost or charge data we are always concerned with the effect of extreme or outlying values on the accuracy of the weight and the stability of the system. In the model outlined extremes may impact both the hospital factor and the APR DRG weight. To balance the effect of extremes in more typical calculations case level charges are often converted to logs after which the mean value is calculated to form the weight (geometric mean). A similar approach is taken for the weight calculation equation (above) with charges converted to their log equivalent. Thus, the equation is redefined as:

$$Ln(C_i) = \omega_n D_{n,i} + \varphi_k J_{k,i} + \varepsilon_i$$

Charges $C_{i,j}$ are obtained from claims and adjusted for trims and edits as appropriate. The log transformation will reduce the impact of outliers. Hospitals enter the regression equation as a series of binary variables as do the APR DRGs. The model is run without constant (the intercept is set to zero as all charges are allocated using the hospital and DRG coefficients) with all DRGs retained. A single hospital is excluded from the binary variable list, while its claims data and binary DRG variables are entered into the regression model, thus enabling the regression to achieve full rank. The regression returns a series of coefficients ω that when exponentiated are representative of the APR DRG **relative** geometric mean charge level of each DRG, and a set of hospital factors ϕ that when exponentiated are representative of the hospital **relative** charge level adjustment required to standardize hospital charge levels to fit the observed pattern of APR DRG charges in the claims data with the least residual. The omitted hospital has no charge standardization factor. Hence its charges serve as the baseline level to which other hospitals are standardized (which in turn informs the absolute charge level of the APR DRG coefficients). At no point is any knowledge of the “correct” level of charge or cost required. Calculations are designed to compute relative values rather than absolute values.

The result is that the hospital factors F, obtained by exponentiating the hospital coefficient ϕ, serve as a series of hospital-specific fractional adjustments that standardize a hospital’s charges to those of the omitted hospital. The exponent W of the DRG coefficient ω is the charge standardized DRG charge estimate that minimizes the squared residual error ε for case i. The use of logs in this calculation decreases the influence of outliers and extreme values while the use of regression focuses upon reduction of squared residual, which encourages fitting to larger values, thus (slightly) acting in opposition. While the result is similar to that provided by the iterative HSRV method, it is neither exactly the same nor intended to be so.

To generate the final weight set, the ω_n coefficients are exponentiated, multiplied by the number of cases in each APR DRG n, and the mean case weighted coefficient value is calculated across all cases. The normalized weight W_n for APR DRG n is calculated as the ratio of the exponent of the coefficient [$\exp(\omega_n)$] to the case weighted mean coefficient. This is equivalent to a relative weight normalized to a 1.0 average where 1.0 is equal to the mean standardized charge per case.

In summary, the log HSRV method provides an analogous result to the HSRV method that avoids iteration. The model removes bias by standardizing case-mix adjusted charges to the level of a single reference hospital. The result increases stability by the use of log normal transformation to charges but recognizes charge dispersion within DRGs to a greater extent by fitting squared residuals.

PPC weights

The calculation of PPC weights builds from the calculation of log HSRV weights. In calculating PPC weights we retain the hospital charge standardization factors F_k calculated previously. The standardization factors F_k are applied to the claims charge data such that they standardize hospital charges relative to all other hospitals ($C_{k,i}$ / F_k).

In computing PPC weights we hypothesize that variation in the standardized charge is a product of interaction between the admission APR DRG A for claim i, and the presence of one, or more, potentially preventable complications P reported for the admission within claim i. As with the log HSRV calculation, the n Admission APR DRGs A are mutually exclusive and may be considered a binary variable that is either present 1 or absent 0 but not both. The c potentially preventable complication categories P are also either present or absent and may be considered binary (1,0) variables. However, multiple PPCs may be reported for the same patient admission i.

In computing the standardized marginal PPC charge, we do not distinguish between preventable and not preventable complications.

The resulting equation is given as:

$$\frac{C_{k,i}}{F_k} = \omega_n A_{n,i} * \Delta_c P_{c,i} + \varepsilon_i$$

Where:

$$A_n = \begin{Bmatrix} 1\ if\ the\ Admission\ APR\ DRG\ of\ case\ i\ is\ APR\ DRG\ n \\ 0\ if\ the\ Admission\ APR\ DRG\ of\ case\ i\ is\ not\ APR\ DRG\ n \end{Bmatrix}$$

$$P_c = \begin{Bmatrix} 1\ if\ the\ PPC\ for\ case\ i\ is\ PPC\ c \\ 0\ if\ the\ PPC\ for\ case\ i\ is\ not\ PPC\ c \end{Bmatrix}$$

The results provide a set of coefficients ω_n for standardized charge estimates of the n admission DRGs A. Coefficients Δ_c estimate standardized charges for the c PPCs P of the marginal increase in standardized charge of a PPC over and above the estimated standardized charge ωn of the admission DRG n for case i. The charge estimates are not log standardized[1]and, by design, are calculated at a level that standardizes to the charge level F of the reference hospital omitted for the log HSRV calculation.

The index value for the log HSRV weights (mean standardized charge per case) serves the same function for the c PPC coefficients Δ_c. The ratio of Δ_c to the case weighted mean coefficient calculated when generating log HSRV weights provides a PPC relative weight W_c for the c PPCs of the same magnitude of the other weights W_n calculated via the log HSRV method for the n discharge APR DRGs.

In summary, PPC weights make use of standardization factors to integrate with the DRG weights.

[1]In weight calculation we are more concerned with central tendency and stability. Complications of care by definition produce outliers making it unreasonable to attempt to mitigate their measured impact.

References

Carter, G. M., & Rogowski, J. (1993). *The Hospital Relative Value Method as an Alternative for Recalibrating DRG Relative Weights.* Retrieved from http://www.rand.org/pubs/monograph_reports/MR156. Also available in print form.

Centers for Medicare & Medicaid Services. Medicare Program; Changes to the Hospital Inpatient Prospective Payment Systems and Fiscal Year 2007 Rates: Changes to DRG Classifications and Relative Weights (2006). Federal Register; Vol. 71, No. 160.

Medicare Payment Advisory Commission. (2005). *Physician-Owned Specialty Hospitals.* Retrieved from http://www.medpac.gov/documents/Mar05_SpecHospitals.pdf.

Notes

1. http://www.medpac.gov/documents/dec05_charge_setting.pdf
2. Grace M. Carter and Jeannette Rogowski, *"The Hospital Relative Value Method as an Alternative for Recalibrating DRG Relative Weights,"* Rand Corporation, 1993, accessed April 15, 2016, http://www.rand.org/pubs/monograph_reports/MR156.
3. Medicare Payment Advisory Commission. (2005). Physician-Owned Specialty Hospitals.
   Mark E. Miller, “Physician-owned Specialty Hospitals” (testimony before the Subcommittee on Federal Financial Management, Government Information, and International Security Committee on Homeland Security and Governmental Affairs, U.S. Senate), May 24, 2005, accessed April 15, 2016, www.hsgac.senate.gov/download/052405miller.
4. "Changes to the Hospital Inpatient Prospective Payment Systems and Fiscal Year 2007 Rates: Changes to DRG Classifications and Relative Weights Centers for Medicare & Medicaid Services," Medicare Program, Federal Register 71, 2006, no. 160.
5. Carter and Rogowski, "The Hospital Relative Value Method."