

Estimation of Variance (2008 Data)

The Hospital Cost Transparency (HCT) project annually reports inpatient payments for APR-DRGs that, during the previous calendar year, had at least 150 observations or at least \$1 million in allowed payments and a minimum of 25 observations. Furthermore, the governing workgroup decided the variance should be estimated after stratifying the data by APR-DRG and severity of illness category (minor, moderate, major, and extreme). This means separately estimating the variance in each cell of a 4xR matrix, with R equal to the number of eligible APR-DRGs and 4 referring to the severity of illness categories.

The difficulty with this approach is that it was known in advance that a substantial number of cells would not have a sufficient number of observations to generate reliable variance estimates. Since the first year's data were ultimately published after combining the severity of illness into two categories (minor/moderate and major/extreme), the 2008 estimated variances used the same combined severity of illness categories.

In the event of small cell sizes, contingencies provided that the analyst should attempt to estimate the variance as follows (in order):

1. Use the pooled variance from the 2008 data if Levene's test fails to reject the null hypothesis that the variance does not differ by severity of illness category
2. The 2008 data supplemented with all observations from the 2007 HCT data set (adjusted for inflation to 2008 dollars)
3. Use the pooled variance from the combined 2007-2008 data if Levene's test fails to reject the null hypothesis that the variance does not differ by severity of illness category
4. The combined 2007-2008 data set supplemented with all observations from the 2006 HCT data set (adjusted for inflation to 2008 dollars)
5. Use the pooled variance from the combined 2006-2008 data if Levene's test fails to reject the null hypothesis that the variance does not differ by severity of illness category
6. The combined 2006-2008 data set supplemented with all observations from the 2005 HCT data set (adjusted for inflation to 2008 dollars)
7. Use the pooled variance from the combined 2005-2008 data if Levene's test fails to reject the null hypothesis that the variance does not differ by severity of illness category
8. Use the variance estimate from 2007 (these methods are described in a separate document)

Preliminary Data Preparation

The 2008 HCT aggregate claims data were first assessed for their suitability for risk-adjustment with the 3M Core Grouping Software. This software assigns a severity of illness category based on diagnoses, procedures, length of stay, patient age, and patient discharge disposition. The risk-adjustment software generally will not provide a severity of illness score if either age or discharge disposition are missing, or if the principal diagnosis is either missing or invalid.

The risk-adjustment can also be influenced by secondary diagnoses and procedures. The standard inpatient discharge record contains nine diagnoses and six procedures. Internal testing revealed that deleting several diagnoses and procedures could alter the assigned severity of illness. Some claims systems do not capture all of the diagnoses and procedures in the discharge record, making risk-adjustment of these records potentially unreliable.

Accordingly, records in the aggregate claims data were required to meet the following risk-adjustment standards:

- Age is not missing
 - This requires the date of birth field to be populated with a valid date
- Length of stay is not missing
 - This requires the admit date and discharge date fields to be populated with valid dates
- Principal diagnosis is not missing
- Claims system captures at least five diagnoses
- Claims system captures at least four procedures

Additional validations were performed to assure the quality of the data and eliminate duplicate records. Details about this extensive process are available in a separate document. The final 2008 HCT aggregate claims file contained 46,656 risk-adjusted and validated records.

Records that did not meet the risk-adjustment standards and records that failed the additional validations were dropped. A separate data file of dropped records was maintained to assess selection bias (see Table 1).

Table 1: Characteristics of Submitted Claims Data

	Risk-adjusted records	Dropped records
Percent routine discharge	91.1%	85.6%
Percent female	61.9%	58.7%
Median age	37	29
Mean length of stay	3.6	6.2
Mean allowed payments	\$13,807	\$12,884

The differences between risk-adjusted records and dropped records are modest, although not negligible. The patients from dropped records tended to be a few years younger, but had longer hospital stays and were less likely to be routinely discharged. The net result of this is patients from dropped records may be somewhat higher risk. The mean allowed payments for these patients were not substantially different. Thus, selecting records for risk adjustment could introduce a small amount of bias favoring patients with slightly less risk, although this does not appear to result in substantial differences in allowed payments.

Inclusion Criteria

Inclusion criteria were then applied to 2008 HCT aggregate claims data:

- Inpatient claims

- Excluding Med-Advantage claims
- Excluding coordination of benefit claims
- Excluding denied claims
- Excluding workers compensation claims
- Patient resided in Oregon
- Patient was treated as an inpatient at an acute-care hospital located in Oregon
 - Excluding VA facilities, Shriner's, and psychiatric hospitals
 - Excluding patients who discharged against medical advice
 - Excluding patients who expired

Estimating Variance

The 2008 HCT aggregate claims data were then stratified by APR-DRG and two combined severity of illness categories (minor/moderate and major/extreme) in order to generate cell sizes. A variance estimation flag was created to identify cells with at least 30 observations and cells where the variance was estimated as the pooled variance. This was repeated with the supplemented data sets: combined 2007-2008 HCT data, combined 2006-2008 HCT data, and combined 2005-2008 HCT data. The variance estimation flags were merged into the 2008 HCT aggregate claims data, the flags indicating the data set from which the variance was estimated. The variance estimate from 2007 was used if cells had fewer than 30 observations in the combined 2005-2008 HCT data and if the pooled variance could not be used (Levene's test rejected the null hypothesis that the variance did not differ by severity of illness category).

With the 2008 HCT data stratified by APR-DRG and two combined severity of illness categories, the standard deviation (SD) of the allowed payments was calculated for each cell with at least 30 observations. The data were aggregated by APR-DRG to calculate the grouped SD of allowed payments. The SD was similarly calculated in the three supplemented data sets. The SD variables from the three supplemented data sets were merged into the 2008 HCT aggregate claims file. Each record in the 2008 HCT data was assigned a SD variable based on the value in the variance estimation flag. Outliers were identified as observations with allowed payments outside +/- 4 SD from the mean. Outliers were removed from the 2008 HCT data and transferred to a data file of excluded records.

The 2008 HCT aggregate claims file was then aggregated by APR-DRG in order to determine the number of observations after removing outliers. Records were again flagged for inclusion based on frequency (APR-DRG had at least 150 observations) or cost (APR-DRG had at least \$1 million in allowed payments and at least 25 observations). Records not flagged were removed and transferred to the data file of excluded records. The data were then stratified by APR-DRG and two combined severity of illness categories (minor/moderate and major/extreme). The allowed payments were summed for each cell and the mean allowed payments, median allowed payments, mean length of stay, and the number of observations were calculated. The results were separately tabulated based on frequency and based on allowed payments.

The data were then further stratified by hospital. Again, the allowed payments were summed for each cell and the mean allowed payments, median allowed payments, and number of observations were calculated. The results were tabulated by APR-DRG for

two combined severity of illness categories (minor/moderate and major/extreme) and each hospital. Cells with fewer than two observations were reported as “0 or 1.”

In order to assess selection bias due to applying the exclusion criteria, several variables (age, gender, length of stay, allowed payments, and discharge disposition) were compared between included and excluded records (note that dropped records are not part of this analysis). In addition, these variables were compared after filtering out records involving pregnancy and childbirth, since these records represent a substantial proportion of the data set and will tend to focus on narrow ranges of age, length of stay, and allowed payments.

Data were analyzed with SAS version 9.1, 3M[®] Core Grouping Software version 2009.2.1, and Microsoft Office[®] Excel 2003.

Results

Originally 121 unique APR-DRGs were identified for further analysis based on either frequency or cost. The resulting 2 x 121 matrix contained 242 cells, although 1 cell was null after generating cell sizes. Of the remaining 241 cells, the variance was directly estimated from the 2008 HCT aggregate claims data for 164 cells (68.0%) and was estimated as the pooled variance for 57 cells (23.7%). The variance was directly estimated from the combined 2007-2008 HCT data for 6 cells (2.5%) and was estimated as the pooled variance for 4 cells (1.7%). The variance was directly estimated from the combined 2006-2008 HCT data for 4 cells (1.7%); the pooled variance was not used as, for all remaining cells, Levene’s test rejected the null hypothesis. The variance was directly estimated from the combined 2005-2008 HCT data for 2 cells (.8%); the pooled variance was not used as, for all remaining cells, Levene’s test rejected the null hypothesis. Finally, the variance was estimated using the 2007 variance estimates for 4 cells (1.7%).

A total of 59 unique APR-DRGs had at least 150 observations. A total of 121 unique APR-DRGs had at least \$1 million in allowed payments and at least 25 observations. This represents a total of 38,818 observations and over \$513 million in allowed payments; 7568 records were excluded from further analysis. Diagnoses and procedures from six major clinical domains (obstetrics/gynecology, childbirth, digestive system, respiratory system, cardiovascular system, and musculoskeletal system) accounted for 88 unique APR-DRGs, over 85% of the observations (33,222) and about 84% of the total allowed payments (\$429 million).

APR-DRG 560 (vaginal delivery) had the highest total allowed payments (\$33,468,445) and APR-DRG 640 (normal newborn) ranks had the largest number of observations (6012). Diagnoses and procedures involving the cardiovascular system accounted for 22 of the 121 APR-DRGs while diagnoses and procedures involving the digestive system accounted for an additional 19 APR-DRGs. Diagnoses and procedures in the obstetrics/gynecology and childbirth domains accounted for 17,747 observations, or approximately 46% of the included observations.

After filtering out records involving pregnancy and childbirth, differences were generally modest when comparing included records to excluded records in the risk-adjusted

aggregate claims file (see Table 2). Included records had a higher percentage of patients routinely discharged and a lower percentage with major/extreme severity of illness, indicating perhaps a slight bias toward lower risk patients. This did not result in substantially lower mean allowed payments; in fact, the mean allowed payments were higher. Overall this indicates that, after filtering out records involving pregnancy and childbirth, the included APR-DRGs do not produce an egregiously biased subset of records from the risk-adjusted aggregate claims data.

Table 2: Characteristics of Risk-Adjusted Claims Data

	Included records	Excluded records
Percent routine discharge	87.5%	81.8%
Percent female	48.3%	49.7%
Median age	54	50
Mean length of stay (days)	4.0	4.1
Minor/moderate severity of illness	82.8%	76.1%
Mean allowed payments	\$18,838	\$16,758

The magnitudes of the differences in median age, percent female, and mean allowed payments are substantially different if including pregnancy and childbirth records. It should be noted that relatively few pregnancy and childbirth records end up being excluded from the risk-adjusted aggregate claims data, so the risk of selection bias is minimal for pregnancy and childbirth records. However these records, since they are a very large subset of the aggregate claims data, can affect the assessment of selection bias in other records if not filtered out.