July 2003

A Primer on Interpreting Hospital Margins

Background

Many rural hospitals, particularly small ones, have historically struggled to survive financially. In order to develop policy regarding Medicare payment to rural hospitals, it is often necessary for policy makers to evaluate current hospital operating profitability. Analysts often utilize hospital margins as a measure of rural hospital financial health, and the gap in average hospital margins between urban and rural facilities is frequently referenced as an indicator of the need for change to the Medicare payment system. In addition, simulations of predicted margins are used as a tool to determine the relative utility of policies aimed at providing relief to struggling rural facilities and to project the potential impact of a proposed policy on future financial performance.

Policy decisions are sometimes made on the basis of average hospital margins, aggregated across the industry or sub-groups of the industry, such as those based on geographic urban or rural location or hospital size. Because there are a large number of rural hospitals, each of which accounts for a relatively small amount of Medicare expenditures, how average hospital margins are calculated can have an impact on the perception of need for legislative and/or regulatory relief.

This Primer answers questions about the most commonly used measures of over-all or payer-specific profitability, total margins and operating margins, and addresses the different ways in which these measures are commonly aggregated when they are used in descriptive studies or regulatory impact statements. The measures are very similar, making it easy for policy analysts to overlook the slight definitional differences when comparing study findings or recommendations from different sources. If careful attention is not paid to the choice of the measures and the method of aggregation, however, there is a risk of misinterpreting differences across groups, or over- or under-interpreting trend data.

Produced by the North Carolina Rural Health Research and Policy Analysis Center,
Cecil G. Sheps Center for Health Services Research, 725 Airport Road, CB#7590
The University of North Carolina at Chapel Hill, Chapel Hill, NC 27599-7590
http://www.shepscenter.unc.edu/research_programs/rural_program

What is the Definition of a Hospital Margin?

A hospital margin is the ratio of hospital profits to hospital revenue. There are two different margins that are frequently used as measures of over-all profitability in health care: (1) total margin, and (2) operating margin. "Total margin" expresses the difference between total revenue and costs as a proportion of total revenue. Included in revenue in the total margin is "non-operating income" — that is, revenue from contributions, public appropriation and other government transfers, investments, and income from subsidiaries or affiliates. When a hospital's margin is computed only with revenues and costs related to patient care, it is usually called an "operating margin", which expresses the difference between operating revenue and costs as a proportion of operating revenue. In most business settings, the numerators in both ratios would be referred to as a "before-tax profit," but in the language of government and non-profit entities, it is referred to simply as "net income," or, more formally, as "the surplus of revenue over expenses".

(1) Total Margins:

$$\frac{\text{total revenue - total cost}}{\text{total revenue}}$$

(2) Operating Margins:

$$\frac{\text{operating revenue - operating cost}}{\text{operating revenue}}$$

For each of these measures respectively, the ratio will be positive if the facility has a total or operating profit, zero if it is at break-even and negative if it has a total or operating loss (it is assumed that a facility does not have negative operating revenue). The ratio is often expressed as a percent: a hospital with total profits of $5 million earned on $100 million of revenue would have a 5% total margin.

How Do Operating and Total Margins Relate to Each Other?

The median values for both operating and total margins for acute care hospitals in the U.S. for the federal reporting periods 1996 through 1999 are shown in Figure 1. Although both margins tend to move in similar directions over time, they each measure slightly different things. For example, median operating margins (the bottom line in Figure 1) were negative in three of the four years, indicating that more than half the hospitals were operating at a loss. However, for these same years, the median total margins were positive (top line), showing that many of the hospitals that lost money on operations were able to make up their losses through other sources of support.

Figure 1: Median Values of Total and Operating Margins Over Time
(as Reported by Participating Medicare Acute-Care Hospitals)

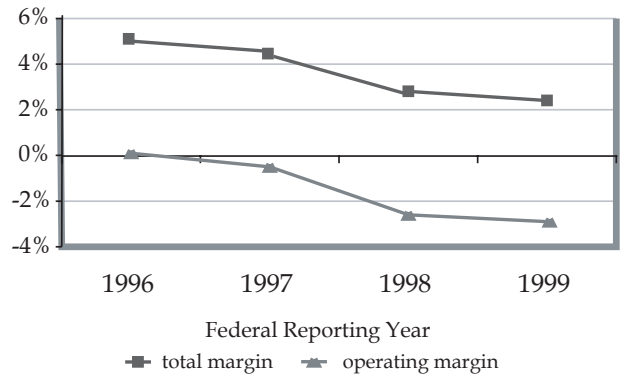

Source: Author's computation from HCRIS reports, FY 1996 – FY 1999.

Which Margin Should be Used to Evaluate Hospital Performance?

While both operating and total margins can be found in policy analyses, and the ratios are similar in construction, they are not interchangeable. One measure may be more appropriate than the other, depending on the objective of the study.
Some public hospitals rely extensively on government transfers that will not factor into the calculation of the operating margin. If public monies are offered only sporadically (for example, in varying amounts that are intended to support a facility through a difficult year), then the operating margin will be a good indicator of that hospital's expected financial performance. If, however, the appropriation is a regular, budgeted item (for example, a stable source of funding to offset indigent care write-offs), then the hospital's price structure probably reflects this other source of revenue. The operating margin, computed without this income, will appear low or even negative, and will give a distorted picture of the expected performance and general financial health of the institution.

What is a Payer-Specific Margin?

In order to assess whether a program is reimbursing hospitals adequately, regulatory analyses and rate studies frequently focus on margins attributable to a specific payer group, such as Medicare or Medicaid. Total margins cannot be meaningfully computed for a single payer group, since by definition, they include non-operating revenue, which cannot come from an insurer. It is possible to compute operating margins based on the payments and estimated costs of patients insured by a specific payer, provided that there is a standardized method for apportioning costs to one patient over another. The Centers for Medicare and Medicaid Services (CMS), the Medicare Payment Advisory Commission (MedPAC) and several other federal and congressional agencies frequently compute Medicare margins in the course of evaluating rates under Medicare's various prospective payment systems (PPS). These margins are based on the data from annual cost reports that are filed by hospitals, nursing homes, clinics and home health agencies. Payer-specific measures are also computed for other non-government programs, but they are less common because it is difficult to estimate payer-specific costs.

Why Do Different Analysts Sometimes Produce Different Estimates of the Average Margin for the Same Group of Hospitals in a Given Year?

Margins are frequently aggregated across hospital groups, or from year to year. Percentile distributions and simple averages are the most common ways of summarizing these data. The 25th, 50th or 75th percentile of a given measure is fairly easy to compute, present and interpret, but averages are not as straightforward. Because margins are ratios of two values, what is labeled as an average may be the simple arithmetic mean of the individual ratios of facilities within the group, or it may be an aggregate measure that is a ratio of the sums of the original information that went into the margin.

Often, the difference between these two types of averages is not properly documented in published tables, but there can be a very important distinction: when we compute a *simple* average of any given profitability ratio, a small facility will have as much influence as a large one on the final average. In contrast, an *aggregate ratio*

(which is computed from the sum of all the information in the numerator divided by the sum of all the information in the denominator) is effectively a weighted average; therefore, the hospitals with larger denominators will have more influence on the resulting summary measure. Figure 2 provides a simple example of this for PPS operating margins, using hypothetical information for four hospitals of varying size with different levels of profit or loss.

Figure 2.

	Hospital 1	Hospital 2	Hospital 3	Hospital 4
PPS Profit	$-5000	$ 0	$80,000	$2,000,000
PPS Payments	100,000	500,000	800,000	10,000,000
PPS Margin	-5.0%	0.0%	10.0%	20.0%

Simple Average Margin $= \frac{-5 + 0 + 10 + 20}{4} = 7.8\%$

	Hospital 1	Hospital 2	Hospital 3	Hospital 4
PPS Profit	$-5000	$ 0	$80,000	$2,000,000
PPS Payments	100,000	500,000	800,000	10,000,000

Aggregate Average Margin $= \frac{-5000 + 0 + 80{,}000 + 2{,}000{,}000}{100{,}000 + 500{,}000 + 800{,}000 + 10{,}000{,}000} = 18.20\%$

In this example, the large hospital has much better profitability than the three smaller ones, so the aggregate ratio is higher than the simple average of the ratios. The two methods of calculating average margins would only result in similar values if small facilities tended to have the same margins as large ones, or if all facilities were the same size. However, among U.S. hospitals, there are many more small facilities than large ones, and these smaller facilities tend to have lower inpatient PPS margins.

How Much Difference Does the Method Used to Calculate Average Margins Make?

The method used to calculate average margins makes quite a bit of difference when comparing urban/rural differences in profitability, as the gap between a hospital group average margin (that is a simple average) and one that is an aggregate ratio average (effectively, a dollar-weighted average) can be substantial. The average difference between urban and rural hospitals based on the aggregate average ratio is much more striking than the difference based on the simple average (Table 1). To the extent that hospitals within the subgroups are more similar in size, there will be less difference between the two averaging methods for any

Table 1: What's the Difference? Comparing Simple Average to Aggregate Average Medicare PPS Operating Margins

	Median (50th Percentile)	Simple Average PPS Margin	Aggregate Average PPS Margin
Rural Hospitals	2.6%	1.6%	3.6%
Urban (MSA) Hospitals	9.7%	8.7%	13.3%
All Hospitals	6.6%	5.6%	12.0%

Source: Author's computations from HCRIS Reports, FY 1999.

given group. However, for urban and rural subgroups, and for all hospitals as a group, the aggregate average inpatient PPS margin has always been higher than the simple average, with the median falling somewhere in between.

Figure 3 provides another example, looking at changes for all acute care hospitals over time. Using the aggregate ratio, margins declined by 21% from 1996 to 1999. Using the simple average across hospitals, the margins declined by 47%.

Figure 3: Inpatient Medicare PPS Operating Margins, by Year

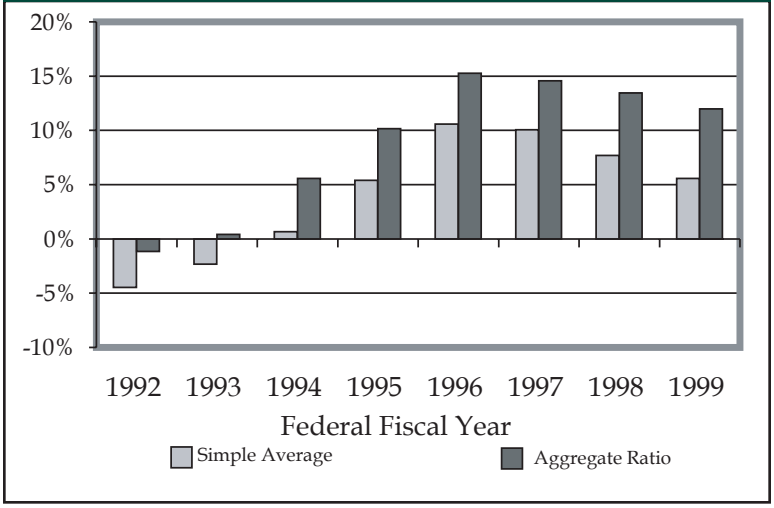

Source: Author's computations from HCRIS Reports, FY 1992 – FY 1999.

Which Average Margin is Better to Use, the Simple or the Aggregate?

Which margin is the right one to use depends on the underlying question to be addressed. The aggregate margin would always be more appropriate for understanding the total budget impact of program changes. MedPAC routinely uses aggregate average ratios to summarize Medicare margins in their annual reports to Congress, and this is the appropriate measure when the policy question of interest is related to overall Medicare spending. A simple average might be better to identify the impact of program changes across individual hospitals. But, neither measure tells us how margins are distributed across hospitals, or what proportion of hospitals is losing money. As was seen in the example in Figure 2, both the aggregate and simple average margins were positive, but only two of the four hypothetical hospitals were actually profitable. When the policy questions relate to concerns about the distribution of margins across facilities, CMS and MedPAC usually report percentile distributions, or they present the aggregated margins by subgroups of hospitals (for example, by teaching status, bed size or urban/rural location).

When Interpreting Hospital Margins, What is Revenue and What is Cost?

Margins computed for individual payer groups warrant close scrutiny for other reasons, having to do with how payer payments and payer costs are defined. Rate analysts generally define payments as the total amount that should be received for the service, assuming that the parts that are the patients' responsibility are collected in full. From a regulatory or rate-setting perspective this is appropriate, because whether coinsurance and deductibles can be collected is a separate policy issue from rate adequacy. By convention, the Medicare margins follow this line of thinking and treat ALL patient balances as if they are or will be fully collected. The convention has merit if the purpose of payer-specific margins is, for example, to evaluate the adequacy of Medicare rates, but not if the purpose is to evaluate the industry's losses on a specific group of patients. Medicare coinsurance and deductible amounts often

have to be written off, and since the Balanced Budget Act of 1997, the Medicare program reimburses providers for only a portion of write-offs attributable to these patient balances.

Definitional problems on the cost side are even more complex. First of all, standard financial accounting reports are designed to identify expenses by type—such as salaries, supplies, or equipment depreciation—but not necessarily by product. Additional cost accounting procedures are necessary for a facility to identify the cost of a particular service that can then be matched to a particular payment. Federally mandated cost reports follow specific rules to allocate administrative and other indirect costs to each patient care service area, and then apportion the allocated service-specific costs to Medicare and Medicaid patients based on the filed claims data. A payer-specific margin computed for other third party payers may not have followed the same cost allocation or apportionment procedures. Comparing ratios from one payer to another may be misleading.

In addition, when evaluating Medicare and Medicaid measures, it is important to know that the programs disallow certain types of costs, such as marketing or professional salaries that exceed regulatory limits, and that these amounts are simply adjusted off the cost reports as if they were not incurred. In most cases, the amounts are not big enough to alter a margin by much, but for some institutions, the effect could be substantial. A more important disallowance that occurs for all hospitals is the set of statutory outpatient cost reductions. These are across-the-board reductions in the computed outpatient costs that have been a part of the regulations governing the outpatient cost-based payments since the late 1980s. When federal agencies compute margins, they usually add back these statutory reductions, but other investigators may not be familiar enough with cost reports to know how to adjust the data.

The statutory reductions on outpatient costs arose from a more complicated problem that should be taken into account whenever Medicare margins are computed for one type of service, but not for all. Medicare has had a prospective payment system in place for inpatient care for nearly 20 years, but it continued to provide some cost-based reimbursement for outpatient services until 1999. Cost analysts generally believe that hospitals had an incentive to maximize reimbursement by allocating more overhead costs to outpatient areas. As a result, inpatient margins are thought to be overstated (as compared to what they would have been in the absence of reimbursement incentives) because correspondingly less overhead may have been allocated to inpatient service areas. The outpatient cost reductions were put in place in part to counteract this, although there is little empirical evidence to support the actual level of the reductions. The statutorily reduced costs were used to compute the average costs per unit of service that serve as a basis for the new rates under outpatient PPS. Thus, artificially low outpatient margins are built into the new payment system.[1]

[1] Using a similar argument about inflated hospital overhead, the new prospective rates for skilled nursing and home health services were set to reflect expected average costs after down-weighting the pool of costs contributed by hospital-based units.

With so much program-to-program distortion designed into the different PPS rates, an accurate measure of hospitals' Medicare profitability can only be obtained if the payments and costs of all of the Medicare services are combined. Recognizing this, MedPAC has begun including an overall Medicare margin in its reports to Congress for the past three years, but the data and procedures for these estimates are more complicated than those for the inpatient margins. There are some Medicare services that are paid on fixed fee schedules (such as outpatient diagnostic lab and durable medical equipment) that never appear on the cost reports, and for which no national data are available. The fee-based payments are well below cost for hospital lab services, and outpatient diagnostic lab services often make up a disproportionately large share of business in small rural facilities. It is likely, therefore, that the overall Medicare margin estimates systematically overstate the ratios for smaller rural hospitals, since services on which they typically lose money (and disproportionately rely on) are not included in the calculation.

What Other Measures Do Policy Analysts Use to Determine Hospital Profitability?

Although not seen as frequently as hospital margins, another payer-specific measure of hospital profitability that is found in the policy literature is the payment-to-cost ratio.

(3) Payment-to-cost ratio (usually payer-specific): $\dfrac{\text{patient payments}}{\text{patient cost}}$

Payment-to-cost ratios are rarely used to analyze overall performance, but are often applied for payer-specific studies. The payment-to-cost ratio is derived from the same information as the operating margin, but it expresses payments relative to cost rather than relative to income. For example, payments of $1.15 million for services costing $1 million would have a payment-to-cost ratio of 1.15, or 115%. The value is greater than one if the facility has a profit, equal to one if it is operating at break-even (because payments equal cost), and between zero and one if it is operating at a loss. It would not normally ever be negative, since neither revenues nor expenses are expected to be negative. If you were to subtract one from the payment ratio, this measure would be equivalent to expressing profit as a proportion of costs, and the scale would be similar to that of the traditional margin. Payment ratios may be somewhat more intuitive to rate and budget analysts, especially those studying health care sectors that were once operated under retrospective cost-based reimbursement, while Margins are the more widely accepted measure in business finance and academic studies.[2]

[2] The choice of expressing profits relative to revenue or profits relative to costs relates to a more basic issue rooted in notions of market-driven versus regulated prices. In micro-economic theory, prices are assumed to be determined in the market place, and so are independent of any individual facility's costs. Thus, to assess a facility's performance, we would want to express profit relative to this independent, or externally determined, measure. In a rate-setting environment, government-administered prices are the item that is being determined, while—in the short run, at least—facilities' costs are treated as the independent item. To evaluate a particular payment scheme, therefore, analysts would compare financial performance under alternative rates, and so would want to consider profits relative to the externally-determined costs.

This work is funded by the Federal Office of Rural Health Policy under cooperative agreement # 6 UIC RH 00027-03.

For more information, contact Kathleen Dalton or Rebecca Slifkin at : (919) 966-5541.

Other recent publications available from the North Carolina Rural Health Research and Policy Analysis Center include:

- Unstable Demand and Cost per Case in Low-Volume Hospitals
- Unpredictable Demand and Low-Volume Hospitals
- Rural Populations and Health Care Providers: A Mapbook
- Design of Enhanced Primary Care Management Programs Operating in Rural Communities: Lessons Learned from Three States

North Carolina Rural Health Research and Policy Analysis Center
Cecil G. Sheps Center for Health Services Research
CB#7590, 725 Airport Road
University of North Carolina at Chapel Hill
Chapel Hill, NC 27599-7590