Hotel-industry Averages

An Inaccurate Tool for Measuring Performance

The commonly used average measures of ADR, RevPAR, and occupancy may be insufficient to see what the “typical” hotel’s performance is really like.

BY CATHY A. ENZ, LINDA CANINA, AND KATE WALSH

Hotel operators and investors use a number of industry statistics as benchmarks to assess current operations and to make forecasts and plans. Three commonly used statistics are occupancy, average daily rate (ADR), and revenue per available room (RevPAR). ADR is the average daily rate per rented room, or the mean price charged for all hotel rooms sold in a given period. RevPAR is calculated by dividing revenue by the number of rooms available for sale. Occupancy is calculated by dividing the number of rooms sold by the number of rooms available and multiplying by 100. Most often summarized as means (averages), those three measures have become common benchmarks by which the lodging industry makes performance comparisons.

By analyzing and interpreting the key industry averages of ADR, RevPAR, and occupancy, it is possible to understand the nature of those averages. In particular, it is worth assessing whether the means of those statistics may mask variations

in the pattern of industry demand and, if so, the extent to which those variations occur. This article reports our analyses of the nature of the industry averages commonly in use to establish performance benchmarks. It also explains the usefulness and accuracy of two other statistical measures, namely, the median and modal values of industry-performance indicators.

Our study is based on monthly demand data for brand-name hotels in the United States for the 13-year period between January 1988 and December 2000, followed by additional analysis on a daily basis for the months of September 2000 and September 2001. The data were drawn from the Smith Travel Research database, which is effectively a census of brand-name hotels in the United States. This comprehensive sample captures over 98 percent of the branded hotel

© 2001, CORNELL UNIVERSITY

inventory. Thus, it is widely considered to be fully representative of all branded hotels in the United States.

How Averages Work

Most people are familiar with the concept of the mean, or average, which gives the value of the central tendency or location of data. The arithmetic mean is a single number that gives the central location of a set of data.¹ The mean is the value obtained when one divides the sum of all values in the data set by the number of items in the set.

A chief problem with the arithmetic mean is its tendency to be distorted by extreme values at either end of a distribution. Indeed, the mean can be so affected by extreme values that it is pulled in the direction of those extremes. As we demonstrate below, this occurs in the hotel industry. If a major market such as New York City or a group of key markets has extremely high ADR figures (as New York and other destinations do), they can pull the entire U.S. ADR upward. In such a situation, the mean may not be representative of the typical ADR in the industry overall, and thus the mean, or average, actually masks the industry's common ADR patterns.

When the arithmetic mean is influenced by extreme values, two other statistics that are measures of central tendency, the median and the mode, may be useful for analysis and comparison.

The median captures the most central or middle value. Put another way, the median is the value below and above which lie an equal number of data points. Unlike the mean, it is the middlemost value in a set of numbers. Since it is based more on the size of the sample than on the numeric values, it has the benefit of being relatively free from the distortion experienced by the mean when a distribution contains extreme values. For instance, the median for ADR tells us that 50 percent of branded hotels are above and

EXHIBIT 1

The performance of branded U.S. hotels between 1988–2000

	Mean ("average")	Median ("middle")	Mode ("typical")	Percentage below the mean
RevPAR	\$41.55	\$36.00	\$30.00	60.98%
ADR	\$63.43	\$56.00	\$47.00	63.30%
Occupancy	63.03%	65.00%	70.00%	47.67%

Total number of observations = 1,817,647

Note: Data were adjusted to year-2000 dollars

50 percent of those hotels are below a particular number. In our study, the average ADR for branded hotels over the 13-year period was \$63.43, while the median for the same period was \$7.43 less, at \$56.00.

The mode shows the value that occurs with the most frequency. Unlike the median and the mean, the mode always appears as a value in the data set. It is often called the most fashionable value (i.e., *à la mode*), because it captures the most typical or representative value located where the data have maximum cluster. Modes for the three performance measures in this study show the industry's most common RevPAR, ADR, and occupancy values. Returning to overall industry ADRs, the U.S. mode is \$47.00, or \$16.43 less than the industry mean. Thus, typical hotels in the U.S. were reporting ADRs below the average during the 13-year period of our study. By including the median and the mode in our analytical toolkit, we can get a fuller and somewhat different picture of the central tendency of lodging performance. With those three measures of central tendency in hand we turn to our study to examine the pattern of frequency distributions in lodging demand since 1988.

A Look at the Means, Medians, and Modes of Lodging Demand

Exhibit 1 shows the overall mean, median, and mode for each of the three key elements of demand—RevPAR, ADR, and occupancy—during the 13-year period. All of the data were adjusted

¹ The average, or arithmetic mean, is the total of the values of a set of observations divided by the number of observations. The calculation is noted as follows $\bar{x} = \sum x_i / n$. For a good discussion of central tendency, see: Morris Hamburg, *Statistical Analysis for Decision Making*, second edition (New York: Harcourt Brace Jovanovich, 1977).

EXHIBIT 2

ADR, RevPAR, and occupancy for U.S. branded hotels, 1988–2000

to year-2000 dollars using yearly consumer-price-index values to control for inflation. We used a total of 1.8 million observations to calculate our statistics.² As the table in Exhibit 1 shows, the industry average RevPAR and ADR are both higher than the typical (mode) or middle (median) points of those hotel-performance measures. At \$41.55, the mean RevPAR is \$11.55 higher than either the median or mode, while the ADR value of \$63.43 is \$16.43 higher than the modal value and \$7.43 higher than the median values.

By looking at the distribution of brand hotels' RevPARs and ADRs we see that 61 percent reported RevPARs below the industry average, and 63 percent recorded ADRs below the industry average. If one relied solely on the average, one might conclude that most hotels enjoy numbers close to the mean, but, in fact, many hotels fail to reach that average. Instead, there's a small number of hotels that are substantially above the average, and those properties distort the mean.

A look at occupancies shows the opposite pattern from that of ADR and RevPAR. The typical hotel occupancy (mode) of 70 percent is higher than the average (mean) of 63 percent. The distribution reveals that 52 percent of U.S. hotels experience higher occupancy than the average value. The extremely low occupancy of some hotels pulls the industry average down to a figure that is 7 percent lower than the typical occupancy levels (mode) and 2 percent lower than the middle point (median). The substantial differences among the average (mean), median, and mode suggest that reliance on just the average could be misleading.

Over the period of time covered by our study, the industry averages for ADR and RevPAR overstated the performance of the typical hotel. Those averages were skewed by a small group of hotels or markets that have much higher ADRs and RevPARs than do all the rest. By the same token, typical hotel occupancies have been understated

² We made our computations of the demand measures monthly at the property level. First we computed each of the demand measures (i.e., RevPAR, ADR, and occupancy) for each hotel per month. We computed the statistics giving equal weight to all hotel properties. Then, we recomputed the statistics while weighting each hotel's performance by its size (i.e., rooms available). Whether the statistics were computed using an equally weighted or rooms-inventory-weighted approach, we found the same pattern of results.

by reliance on an average that is distorted by extremely low occupancies of a small group of hotels or markets.

Distortion in Lodging Demand—Skewness

If the mean, median, and mode for the hotel industry's ADRs, RevPARs, and occupancies were all the same value or close to the same number we would be able to conclude that the hotel industry experienced what is known as a normal, or symmetrical bell-shaped distribution for those statistics. That did not occur with our hotel-industry data, however. Instead, lodging demand displays a non-symmetrical, or skewed, distribution.

This skewness is a key factor in producing the inconsistencies that we have observed with regard to lodging-performance benchmarks. Skewness results when the frequency distribution has a heavy mass of extreme values. A symptom of skewness occurs when extreme values pull the mean away from the median and the mode in the direction of the extreme values. The median is also pulled away from the mode in the same direction as the average, but the median is not affected as much by extremes as is the mean.

This skewed pattern of lodging demand is shown in Exhibit 2. Plotting the frequency distribution of hotels' ADR, RevPAR, and occupancy numbers for each month over the 13-year period reveals the shape of the distribution. The mode is the highest point on the curve, which is the area where the largest number of hotel values cluster.

When the mean (average) is less than the median (middle), as it is for occupancy, the distribution is said to have a negative skewness (i.e., skewed to the left). In contrast, for RevPAR and ADR the mean is greater than the median and mode, and hence the distribution has a positive skewness (i.e., skewed to the right). The position of the mean, median, and mode in these graphs reveals the distance between those statistics and shows the asymmetry in U.S. lodging demand.

Given the above results for overall demand, hospitality managers may have one of two reactions. The first is that, considering the differences in the values of the mean, median, and the mode (in addition to the wide variability of those

demand statistics), using overall U.S. lodging-demand averages alone may prove to be too simplistic. The second is that they already ignore overall averages and focus their analysis on their own local market's lodging-demand data. While analysts and chain-wide brand-pricing strategies may be focused on overall industry measures, many property-level managers and consultants doing feasibility studies use their own markets and competitive sets to set performance benchmarks. Nevertheless, decision makers may not

Over the 13 years covered by our study, the industry averages for ADR and RevPAR overstated the performance of the typical hotel.

fully grasp the variability within markets. As such, it is important to understand the patterns of differences that may be due to market location or segment. To refine our understanding of the patterns of lodging demand we now turn to an exploration of key markets and market segments.

The Differences in Major Markets

Our analysis revealed that substantial variation exists in overall lodging demand. To determine whether this variation is the result of distinctive patterns in specific markets, we explored whether primary markets (which we defined as the top-25 markets) showed significantly different demand patterns from the remaining markets. By conducting a series of statistical tests to determine whether the arithmetic means for the primary markets were similar to those of the remaining markets, we found that the major markets have significantly different performance levels. By splitting the branded hotels into two groups, one composed of the top-25 markets and the other consisting of all other U.S. markets, we clearly see the difference in their demand patterns (Exhibit 3, on the next page).³

³ Using a series of *t*-tests, we tested for differences between the top-25 markets and all others on ADR, occupancy, and RevPAR. To see a more detailed paper reporting the results of our statistical tests, please contact us.

EXHIBIT 3

Performance of branded hotels in the top-25 markets compared to all others, 1988–2000

	Mean ("average")	Median ("middle")	Mode ("typical")	Percentage below the mean
RevPAR				
RevPAR for top-25 markets	\$52.56	\$45	\$30	60.30
RevPAR for all other markets	\$37.73	\$33	\$26	58.99
ADR				
ADR for top-25 markets	\$76.35	\$67	\$52	61.57
ADR for all other markets	\$58.93	\$54	\$47	60.54
Occupancy				
Occ. for top-25 markets	66.36%	69%	75%	45.73%
Occ. for all other markets	61.87%	63%	70%	46.64%

Number of observations for top-25 markets = 469,299

Number of observations for all other markets = 1, 348,348

The findings reported in Exhibit 3 show that the average RevPARs and ADRs in both primary and other markets are higher than the typical hotels' performance levels in those markets. The degree of difference between the average and the typical hotel was substantially greater for the top-25 markets than for the other hotels. In the top-25 markets, the typical hotel's RevPAR of \$30.00 was \$22.56 below the average RevPAR. In contrast, the typical RevPAR of \$26.00 in other markets was only \$11.73 lower than the average. The big gap between typical (mode) and average (mean) hotels for the top-25 markets suggests that the average is unduly influenced by the extremely high RevPARs of a relatively small set of hotels. The smaller gap between the typical and average hotels in all other markets suggests that the bulk of U.S. markets have more normal distributions without extraordinarily high RevPARs.

Our exploration of ADRs revealed a similar pattern. For the top-25 markets the gap between the typical hotel and the average hotel was \$24.35, while the gap was \$11.93 for ADRs in other markets. All other markets were less volatile than the top-25 markets, indicating a narrower range of demand values. ADRs for some markets in the top-25 are so high that they again unduly distort the average and push it to a higher level than the typical hotel's ADR. In both the case of ADRs and RevPARs the average overstates performance when compared to the typical hotels.

Average occupancy figures were not so dramatically different between primary and other markets. While occupancies were 4.5 percent higher in the primary markets when compared to all other markets, the averages were consistently lower than the typical hotel occupancies in both types of markets. The primary markets' average occupancy of 66.26 percent was 8.64 percent lower than the modal hotels, while the average occupancy in all other markets (at 61.87 percent) was 8.13 percent lower than the typical hotels in those markets. Again, the results reveal negative skew such that the average occupancy is lower than the typical hotel occupancy. So, the average understates the number of rooms that are efficiently filled with guests. (Our market comparisons showed significant differences between primary and other markets, but those oc-

cupancy differences were smaller than were the differences for RevPAR and ADR.)

In addition, the elevated levels of demand in some markets may overstate demand for secondary markets (when using averages), but that may also be true for the top-25 market itself. The degree of variation in demand was found to be so great in the top-25 markets that treating them as though they are a group may be problematic. Since top-25 markets have great variability in their performance, we next turned to an examination of key cities within the top-25 markets to more fully understand the degree to which variability in demand patterns exists.

RevPAR in Ten Key Cities

While our study showed that the top-25 markets were distinct from secondary markets, the high volatility and skewness in demand would suggest that key cities should be examined separately to see how similar they are to each other. We chose ten key cities for this analysis: Atlanta, Boston, Chicago, Los Angeles, Miami, New York, Philadelphia, Phoenix, San Francisco, and Washington. We chose RevPAR as the performance indicator for our tests because it is often considered the most critical measure of operating performance, and by definition, encompasses an element of rate and room supply.⁴ Conducting a series of statistical tests, we explored the differences between the mean RevPAR of each of the ten cities compared with the mean of the top-25 markets excluding that city. Based on this analysis we found each of the ten key cities to have RevPARs significantly different from the RevPARs of the remaining top-24 markets. Thus, each top-ten city is distinctive from the remaining top markets.

Exhibit 4 shows the means, medians, and modes of RevPAR for each of the ten cities. Those key cities are listed in order from New York (with the largest gap between the mean and mode values) to Miami (with the smallest gap). In comparing a particular key city to the remaining top-24 markets, we found that New York City has the highest degree of variability, as well as the

EXHIBIT 4

Differences in RevPAR for branded hotels in key cities, 1988–2000

Key city:	Mean ("average")	Median ("middle")	Mode ("typical")	Percentage below the mean
1. New York	\$123.10	\$106	\$75	62.07%
2. San Francisco	\$ 76.73	\$70	\$48	57.35%
3. Phoenix	\$ 51.22	\$40	\$27	64.47%
4. Los Angeles	\$ 51.95	\$44	\$30	59.67%
5. Boston	\$ 71.30	\$62	\$55	60.16%
6. Washington	\$ 61.99	\$54	\$46	58.46%
7. Atlanta	\$ 43.91	\$38	\$28	58.43%
8. Chicago	\$ 57.21	\$51	\$42	59.43%
9. Philadelphia	\$ 56.71	\$53	\$44	56.12%
10. Miami	\$ 59.50	\$53	\$49	60.21%

Note: The difference between the mean RevPAR of each city compared with the mean RevPAR of the top-25 market (excluding that city) is statistically significant. The difference between the distributions of the RevPAR of each city compared with the distribution of the RevPAR of the top-25 markets (excluding that city) is also statistically significant. The cities are presented in order of the difference between their mean and mode from the largest difference to the smallest.

⁴ For clarity we present only the results for RevPAR. The results for ADR and occupancy are available from us upon request.

EXHIBIT 5

RevPAR for key U.S. cities between 1988–2000

largest gap between the average of all hotels and the modal value that represents the typical hotel. (That difference was \$48.10.) San Francisco, Phoenix, Los Angeles, and Boston are the cities with the next-highest degrees of difference between the average and the typical hotel (see Exhibit 4). Interestingly, Los Angeles, Phoenix, and Atlanta have the lowest average RevPARs, while Miami has the smallest gap between the average and the typical hotel RevPAR levels.

The frequency distributions of RevPAR for those top-ten cities are shown in Exhibit 5. New York City has the highest degree of variability in RevPAR, while Atlanta has the lowest. In fact, the RevPAR pattern for NYC is unlike that of any other city—owing both to its wide range of values and its extremely high values. The RevPAR patterns of Boston, Chicago, Philadelphia, San Francisco, and Washington are reasonably similar to each other. It is interesting to note that we have heard Washington being discussed as though its demand pattern were like that of New York City, when in fact the D.C. market's demand patterns are closer to those of Chicago and Philadelphia. Even given the similarities of the five markets we just discussed, the statistical tests and graph of distributions reveal that RevPAR demand in the key cities is not comparable. Moreover, NYC's RevPAR inflates the average substantially for the top-25 markets.

Variability by Segment

Another way to examine the variability in demand patterns is to examine the differences by price segment—to see whether demand in some segments is skewed and to determine the extent to which demand patterns vary from segment to segment. To investigate that matter we looked at the means, medians, and modes for the following five price segments: luxury, upscale, midprice, economy, and budget. In our analysis of price segments, presented in Exhibit 6, we see that the same pattern exists in each segment (similar to the overall industry). That is, for each segment the mean exceeds both the median and the mode for both RevPAR and ADR, while the reverse relationship holds for occupancy. Occupancy is negatively skewed, given that the mean is lower than the median and the mode. While the typical hotel has lower RevPAR and ADR figures than

the average, the differences in those statistics are greatest for luxury hotels, with a \$22.34 difference between the average ADR and the typical hotel ADR. Upscale hotels have only a \$12.59 difference, with midprice, economy, and budget segments reporting differences of between \$7.69 and \$3.51 when comparing the average with the typical hotel in their price segments.

While luxury hotels have the greatest degree of variation in ADRs and RevPARs, this segment has less variation in occupancy level than do the other segments. Budget and economy properties have nearly normal distributions for ADR and RevPAR (little variation and skew in their performance patterns). Thus, industry averages are more useful for individuals who focus on those two price segments. Clearly, caution is warranted for those who want to understand performance in upscale hotels. To a greater extent than other segments, care is necessary when interpreting performance based on averages for luxury hotels.

Occupancies in the Wake of Crisis

On September 11, 2001, the United States experienced terrorist attacks on the cities of New York and Washington that profoundly changed the existing economic activity in the lodging industry. To better understand the actual effect of the events of September 11, we looked at occupancy averages for three different time periods in the month of September. The first period, September 1 through 10, includes the Labor Day holiday, which historically produces low ADRs and occupancies. (In fact, the lowest occupancy for September was reported on September 3.) The second period was September 11 through 16, the week after the terrorist attacks, during which time the fact that some people were stranded in key cities may have artificially inflated occupancy. The last period was September 17 through 29, during which the wide perception was that people were “not traveling.” For this analysis we relied on September’s daily occupancy data for the years 2000 and 2001. By calculating the means, medians, and modes for all of the United States and for New York City and three other key cities most directly affected by the attack (Boston, San Francisco, and Washington), we found a picture of huge variability by looking at both the average and the typical hotel.

EXHIBIT 6

Performance of U.S. branded hotels by segment, 1988–2000

	Mean ("average")	Median ("middle")	Mode ("typical")	Percentage below the mean
RevPAR				
Luxury	\$81.90	\$74.00	\$69.00	60.28%
Upscale	50.34	47.00	43.00	57.05%
Midprice	38.00	36.00	33.00	57.21%
Economy	28.90	27.00	23.00	54.16%
Budget	23.53	22.00	19.00	54.64%
ADR				
Luxury	\$114.34	\$103.00	\$92.00	64.45%
Upscale	74.59	70.00	62.00	59.44%
Midprice	59.56	56.00	52.00	59.95%
Economy	47.51	46.00	44.00	58.58%
Budget	39.21	38.00	34.00	58.04%
Occupancy				
Luxury	70.17%	73.00%	78.00%	44.25%
Upscale	65.54	68.00	75.00	45.16%
Midprice	62.26	63.00	67.00	48.21%
Economy	59.82	60.00	60.00	48.83%
Budget	59.80	60.00	60.00	48.97%

As the performance measures in Exhibit 7 show, New York City and the three other key cities saw substantial declines in occupancy. Before the events of September 11, New York occupancies were 7.42 percent lower than the banner year of 2000. Immediately after the tragedy, year-to-year occupancies were down by 20.12 percent between September 11 and 16, and went further down (by 22.33 percent) in the third period (September 17 to 29). Drops in occupancy for the other three key cities were even larger than in New York City. Looking at the United States overall, we found a 2.79 percent occupancy drop in 2001 from 2000 before September 11, with year-to-year drops of 10.10 percent in the week after the attack, and 9.27 percent in the final

EXHIBIT 7

Occupancy percentages of branded hotels for
September 2001 and 2000

Time Period	Mean ("average")		Median ("middle")		Mode ("typical")	
	2001	2000	2001	2000	2001	2000
September 1–10						
Total U.S.	56.06	58.85	55	58	50	98
Key Cities	64.68	73.53	65	78	98	98
Key Cities w/o New York	63.28	72.29	63	76	57	98
New York	77.02	84.44	79	90	98	97
September 11–16						
Total U.S.	55.73	65.83	55	67	50	98
Key Cities	56.93	84.67	57	92	52	98
Key Cities w/o New York	55.59	84.15	55	91	50	98
New York	69.33	89.45	70	95	91	99
September 17–29						
Total U.S.	56.42	65.69	55	67	50	98
Key Cities	55.09	84.88	53	92	50	98
Key Cities w/o New York	53.79	84.40	52	92	50	98
New York	67.00	89.33	68	95	59	98

Our results suggest that figures for key cities dominate and distort average lodging demand for brand hotels.

September period. While those figures are strong indicators of a decline, the modal or typical hotel's performance shows a far more dramatic drop than the averages reveal.

Exhibit 8 shows the distribution of occupancy levels for the three time periods. What is notable about these figures is that they show that the typical hotel's occupancy, which had been at the 90-percent level, was cut nearly in half, to around 50 percent. The dramatic shift in the typical hotel's occupancy, rather than the average occupancy, has led many in the field to express alarm. On the other hand, those looking at overall averages may conclude that a far less dramatic crisis has occurred. Our statistics and figures show how important it is to consider more than one statistic when attempting to make forecasts of economic effect and future recovery.

Taking All Statistics into Account

The fundamental message our study offers is to proceed with caution when using industry aver-

ages for forecasting and making decisions. Unstable and turbulent environments can produce extreme values that pull the arithmetic mean in one direction or another. As our study has shown, that average alone is not adequate to describe lodging demand fully, particularly since the industry reflects such large variation in markets, key cities, price segments, and environmental (external) factors. Reliance on the average can lead managers to both overstate ADR and RevPAR goals and understate occupancy goals. We now know that the average overstates ADR and RevPAR because of extremely high numbers in some segments (particularly, luxury) and certain markets (notably, New York City). Occupancies are actually higher for the typical hotel than the averages would suggest.

Actions to Improve Decision Making

To ensure that one uses the appropriate demand data to compare against operating performance and to prepare forecasts, we advise augmenting

EXHIBIT 8

The total U.S. room occupancy frequency distribution for September 2000 and 2001

aggregate industry measures with additional statistics, such as the median and mode. Another statistic, the standard deviation, which expresses the extent to which the data are dispersed, can also provide useful information, although we did not discuss it here. Using all those statistics may enable decision makers to more fully explore the possible role of skewness in lodging demand and identify areas where current measures may be distorted.

Our results suggest that key cities dominate and distort lodging-demand averages for brand hotels. Overall lodging data mask huge variability that exists by market and segment. In addition, performance patterns vary substantially within markets and segments. A look at daily

demand in light of the events of September 2001 further supports the observation that the lodging industry has patterns of demand in key cities that shape our impressions of the entire U.S. market.

It makes sense to add measures of performance to the analysis, the better to develop a fuller understanding and more appropriate comparison points for meaningful benchmarking. Such additions would enable managers and investors to conduct more effective analyses and make more precise decisions. In addition, selecting the right comparison or reference city, segment, or market for a benchmark will go a long way toward improving the hospitality decision maker's toolkit. ■

Pictured clockwise from top left:

Cathy A.ENZ, Ph.D., is the Louis G. Schaeneman, Jr., Professor of Innovation and Dynamic Management at the School of Hotel Administration and executive director of the Center for Hospitality Research at Cornell University (CAE4@cornell.edu), where **Linda Canina**, Ph.D., is an associate professor of finance (LC29@cornell.edu) and **Kate Walsh**, Ph.D., is an assistant professor of management (kmw33@cornell.edu). The authors point out that they were equal contributors to this research project and that an earlier version appeared as a Center for Hospitality Research Report. The authors also gratefully acknowledge the assistance of the Center for Hospitality Research and Smith Travel Research in conducting this study.

© 2001, Cornell University; refereed article: submitted on October 29, 2001; accepted with minor revisions on December 10, 2001.