2002

Performance Monitor

Montgomery Blair
Dollar Thrifty Automotive Group

Chris K. Anderson
Cornell University School of Hotel Administration, cka9@cornell.edu

Performance Monitor

Abstract
Revenue management has experienced tremendous growth since its inception in the airline industry. The practice of dynamically pricing a perishable product across different market segments continues to be applied across an ever-increasing set of business arenas. The proliferation of revenue management can be attributed to its microeconomic foundations (grounded in supply and demand), as well as the numerous well-cited success stories (American Airlines, National Car Rental) and failures (People's Express) that follow from its application or lack thereof. The continued success of revenue management hinges upon the ability to link organizational performance to the pricing and capacity decisions of revenue management systems. This link both reinforces the financial gains attributable to revenue management and indicates opportunities for future improvement. This paper outlines Performance Monitor, a phased approach to performance measurement designed and implemented at Dollar Car Rental. We also present and discuss some examples of Phase I, which is currently in use.

Comments

Performance Monitor

Montgomery Blair, PROS Revenue Management
Chris K. Anderson, University of Western Ontario

Revenue management has experienced tremendous growth since its inception in the airline industry. The practice of dynamically pricing a perishable product across different market segments continues to be applied across an ever-increasing set of business arenas. The proliferation of revenue management can be attributed to its microeconomic foundations (grounded in supply and demand), as well as the numerous well-cited success stories (American Airlines, National Car Rental) and failures (People's Express) that follow from its application or lack thereof. The continued success of revenue management hinges upon the ability to link organizational performance to the pricing and capacity decisions of revenue management systems. This link both reinforces the financial gains attributable to revenue management and indicates opportunities for future improvement. This paper outlines Performance Monitor, a phased approach to performance measurement designed and implemented at Dollar Car Rental. We also present and discuss some examples of Phase I, which is currently in use.

Introduction

For good reason, tremendous effort and resources are currently being invested in revenue management systems. Revenue management has a successful history in the car rental industry; in fact it is credited for pulling National Car Rental from the brink of bankruptcy (Geraghty and Johnson, 1997). At the heart of most systems is a science-based model that has a well-defined objective. The first generation of systems typically had a straightforward objective function such as maximizing profit or revenue. Within these systems, many inputs are exogenous and several assumptions are made. The desire to develop improved systems, coupled with the very significant investment required for their development, necessitates a clear understanding of their performance. The system developers (vendors), the organizations that use them, as well as the analysts who implement the system recommendations, all require a clear understanding of system operations and its impacts upon organizational performance. This paper illustrates some of the difficulties associated with assessing the

impact of revenue management systems, and then presents a system currently is use at Dollar Rent A Car, which provides a first step towards achieving revenue management accountability.

Revenue Management Performance

An objective assessment of business performance as impacted by revenue management activities is often difficult. There are three primary difficulties associated with analyzing the performance:

— the nature of an optimization

— lack of a control, and

— a definition of 'good'.

The Nature of an Optimization

Most optimization processes simultaneously adjust variables within defined constraints to maximize some objective function. The magnitude of potential combinations is in direct correlation with the number of factors within the model. To have some degree of reality, many of the complex systems utilize a large number of inputs and thus necessitate the most advanced computer systems. It may be inconceivable to dissect the individual rational within the trade-offs of a complex optimization routine. With this, the 'users' are left with the solutions, the process definitions, and the application of such results.

Lack of a Control

With most scientific experiments, it is paramount to use a base control to compare against. In the real world of revenue management applications, there is no control. No matter how complex a system may be, the methods are still far from incorporating all potential variations in the business environment. Time is of the essence in revenue management, and most industries have an interest in revenue management systems owing to the potential for spoilage. The concept of time alone makes a dual environment — or control impossible. The second significant ingredient that eliminates the possibility for good control is the concept of price elasticity. The concept surrounding the quantification of the sensitivies of price and quantity not only abolishes the ability to formulate a controlled environment, but is also enormously critical within the systems themselves. Although elasticity and cross-elasticity have been measured in similar business settings (Oum, 1989), the range within which these estimates (and others) fall is astounding (Dargay, 1993). There exists an inherent stochastic nature

within elasticity itself that should not be over-simplified. At the heart of the price elasticity is human nature, the freedom of choice, and consumer preference. People can change their minds, act only as rationally as they desire, possess imperfect information and exhibit different individual reservation prices. A single individual can vary his or her own price sensitivity, ceteris paribus. Some may believe that prices will fall before expiration, while others may not. Some may have different inconvenience 'costs' for shopping. This culminates in an inability to develop a credible controlled norm.

In the car rental industry, a revenue management system may change 100,000 rates each day for a 90-day horizon. For each arrival date, location, car class, length of rent, specific distribution channel, a rate code could be altered independently numerous times. It would be absolutely erroneous to make bold assumptions on that which would have occurred had one or more changes not been performed. Because there are thousands of combinations for an infinite number of prices, control comparisons are neither possible nor valid.

A Definition of 'Good'

A soccer team scores three goals. How well did they perform? Is it only measured against their competitor? Did they win or lose? Even if victorious, perhaps the other team was extremely bad. Could they have done better? What is their team's record? How well did they do last year? What if their particular league is one of the premier organizations in the world? One should be able to gain an appreciation for the difficulty in just understanding how quantifying performance measures can still fail to define performance absolutely. With companies, the ultimate measure may truly be profit, and with a soccer team it may be the ultimate final standings. There are numerous difficulties with this single-minded approach. For example, what if the team experiences an injury, or another stronger team enters the league, or vice versa? In the soccer illustration, a player may step in for the injured and perform better. For the company, what if a competitor drags down the market, or what if the industry's demand shifts? Within the company a system's component may improve performance even in light of the final profit numbers. The bottom line is that performance measurement may not be well defined. Even with quantification and benchmarks, there are still perspectives and various levels of aggregations.

The goal of most companies is rather simple and straightforward — maximize profit. So the easiest performance-measuring device would be to measure profit. With most publicly traded companies, this process is done rather efficiently via the equity markets. The problem is determining the causality associated with the profit performance. Specifically, deriving the relationship between applying revenue management systems and the bottom line.

Many developers of large-scale revenue management systems have professed a bottom-line increase of 2–8 per cent or larger (Manugistics, 2001; PROS, 2001). Is this due to the implementation of RMS, or is it more a mere coincidence of a booming economy? As the economy has currently slowed its robust growth, and many companies are experiencing diminishing or even negative profits, are the revenue management systems held accountable? What are these same developers going to use to benchmark and quantify the benefits derived from utilizing a more systematic approach to new businesses?

American Airlines (AA) is one of the few practitioners of revenue management that has published an account of its attempts at monitoring revenue management performance (Smith et al., 1992). American Airlines quantifies revenue gains from overbooking by determining the percentage of available revenues it captures (the difference between no overbooking and perfect overbooking — determined by adding in spoilage and over-sales costs and removing lowest revenue customers on full planes). American Airlines attempts a similar analysis with discount allocation, but indicates the lack of effectiveness due to the nature of the marketplace. American Airlines has gone on to attribute over $1bn in annual incremental revenue to its revenue management systems (Cook, 1998).

Dollar Rent a Car

In November 1997, Daimler Chrysler Motors Corporation divested itself of Dollar Thrifty Automotive Group (DTAG). Currently, DTAG owns Dollar Rent A Car Systems, Inc., and Thrifty Rent-A-Car System, Inc., which operate separate rental businesses and license independent franchisees to rent vehicles under their brands. Based in Tulsa, Oklahoma, Dollar Rent A Car is one of the most profitable in the industry. Over the past five years, it has grown to approximately 460 worldwide locations in about 26 different countries, including 260 in the USA. The fleet size is approximately 210,000 vehicles of which nearly 97,000 are in the USA. Their primary focus is on the price sensitive leisure traveler in the USA. As displayed in Figure 1, DTAG's revenues have grown over the past five years and are nearing $1bn per year.

Dollar can also boast that they have become one of, if not the, most profitable car rental companies in the world. This is due in large part to the revenue management practices that they undertake. In 1998, they began an implementation of a Talus Rental Car Revenue Management System. They have spent significant time and effort in not only utilizing this system, but also in developing several other revenue management tools and practices. The RM department balances a conservative approach with that of strategically capitalizing on science-based systematic approaches to pricing. The

dynamic pricing capabilities at Dollar can be appreciated by its ability to change nearly 100,000 rates per day. This accomplishes two primary tactical achievements. First and foremost, these changes are calculated based on an optimization of forecasted demand, realized variances and competitive actions. At all times the price changes have Dollar's profit interest at heart. Secondly, it diminishes the ability for the competition to react. If a competing company has an objective based on an internal need to build more or less reservations, they will have difficulty formulating an efficient methodology to achieve it. If, for example, they are priced higher than Dollar and thus are not getting as many reservations as needed, they cannot be certain as to an appropriate price at the time due to Dollars' extremely dynamic pricing. Thus they are likely to over- or under-react. Also, the level of detail is greater with Dollar's dynamic pricing than that of the competition. Any competitive pricing move will surely be an overgeneralization for sets of dates or car classes. This 'revenue management' practice of Dollar, coupled with low-cost-efficient operations, could be a reason Dollar surpasses all others in profit (see Figure 1).

The question is asked, however, 'How well is Dollar doing'? This simple and brief background into Dollar highlighted the essence of the performance measurement process. Although the bottom-line profit is the ultimate performance measurement, determining the components and causality between the systems and final results dictates advancements in this arena. From the very fine detail of analyzing the value of each pricing decision, the quality of the system that created it, to all individuals with decision-making abilities, performance measurement is critical.

There clearly needs to be a more direct analysis of revenue management system performance. The initial challenge for Dollar was to answer two performance questions; which locations are performing well, and how well?

Figure 1: Sales and profitability growth at Dollar (000s)

Net Sales
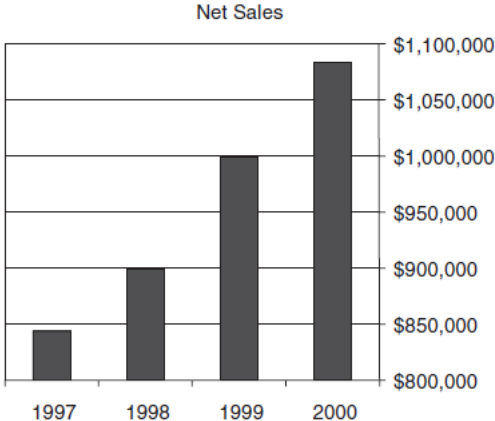
A Performance Monitor at Dollar Rent a Car

The purpose of Performance Monitor is to aid in the economic evaluation of a locations results. Currently, common benchmarks are used in the car rental industry, such as utilization and revenue per day (RPD). Alone, these can be misleading and very subjective. Even superior measurements such as revenue per unit (RPU) fail to capture the entire picture, understand the opportunity cost and address accountability.

This particular system is being developed in several phases to address these issues. Phase I, which is in production and updated daily, calculates objective results and signals opportunities and missed actions. Phase II, currently in development, measures performance relative to potential performance. It quantifies the opportunity cost associated with the business that was not consummated. The assumptions are limited to only those surrounding the interpretation of the quantified results and not so much the results themselves. In other words, it pushes the assumptions further down the performance measurement rung. Both of these two initial phases internalize minimal assumptions, thus maximizing objectivity, and aid in the performance analysis of the entire system.

Future phases will include more quantitative financial numbers. The thought process is to develop an objective system that will ultimately show financial results, quantified opportunity costs, as well as point towards accountability and responsible parities. This in turn may be the foundation for a systematic approach to profit maximization through optimizing comprehensive sound objectives.

Phase I

Specifically, this performance monitor is designed with a basic objective foundation that determines relative performance from multiple perspectives, including daily, weekly and monthly. Since the objective measurements are 'relative', the final results for a single location vary with the group to which it is included. For example, a 'report' that analyses a location's performance relative to the other locations within the same region may post different results from when it is compared with all locations or just itself. Through the use of various perspectives, issues of seasonality, geographic variations and exogenous market conditions are addressed.

The metrics of Phase I are utilization and turndowns; utilization being the fraction of available inventory on rent, turndowns resulting from denying a particular service. Turndowns are broken down into rate turndowns (RTD) and availability turndowns (ATD). RTD reflect the business interested but not willing to pay the current price. ATD reflect an apparent lack of capacity. Generally speaking, high

utilization is reflective of good performance. Large numbers of ATD indicate a potential loss of last-minute, high-value demand. The impact of rate turndowns depends on utilization. A high degree of rate yielding is detrimental in the case of low utilization whereas, a low number of RTD may be poor under instances of high utilization, and inconclusive or indifferent otherwise.

Phase I is a comparative tool. Instead of qualifying exactly what levels of utilization and turndowns are strong/poor (elements of Phase II), relative measures are used, resulting in performance on utilization and turndowns judged between Dollar locations. Utilization and turndown (normalized by location fleet size) averages are calculated as benchmarks. These benchmarks can be for specific locations, regions or all locations and determined for specific days or aggregated by weeks or month. A location can then be compared with the benchmarks to assess when it has exceeded the average (High) or been a sub-par performer (Low). The performance across utilization, ATD and RTD can then be combined to assess overall location performance relative to the benchmarks. Table 1 displays the quality grid. The grid indicates the most preferential mix of utilization and turndowns at the top, which receive a score of 8, the least preferential scoring a 1.

Table 1: Performance quality grid

Utilisation		ATD	RTD	Comments	Score	
High	Top Tier	Low	High	Yielded nicely	8	Better
High		Low	Low	Could have yielded more	7	
High	Mid Tier	High	High	Could have definitely yielded	6	
High		High	Low	Turned away high value	5	
Low		Low	Low	Indifferent — too much fleet?	4	
Low	Bottom Tier	Low	High	Yielded — cost rentals	3	
Low		High	Low	Turned away business	2	
Low		High	High	Very poor	1	

Top tier locations do an excellent job of achieving high utilization while turning away low-value traffic. The Mid tier locations, either attained desirable utilization levels but perhaps did so at acceptance of low-yield customers, or they had poor utilization but may not have been able to do much about it. Bottom tier locations performed poorly with regard to utilization and may have simultaneously turned away potential high-value customers.

The use of relative benchmarks versus quantitative targets allows for the control of environmental factors (e.g. economic conditions, seasonality or competition) as locations can be compared with others in similar regions under similar competitive environments. At the same time, a

location may be compared with its historical performance to gauge whether or not, relative to its past performance, it has been improving. Figure 2 displays the nine different levels of benchmarking, which are part of Performance Monitor. There are three levels of comparison — same location (L), region (R) or all (A) locations — with each level having comparisons done on the day (D), week (W) or monthly (M) level.

Example

Table 2 displays sample information on utilization and turndowns for six locations. The performance at each location is compared with the average of the six, and judged to be higher or lower. Table 3 summarizes the performance across all three metrics and scores the six locations based upon the quality grid (Table 1). Figure 2, output from Performance Monitor, displays how these scores look over a period of time. The patterns that start to develop lead to insights about location performance. When locations remain in a 'tier' of the quality grid for consecutive days or weeks, they become indicators of certain conditions. For example, a group of 4s might indicate an oversized fleet that might be better served at another location, or a series of 1s and 2s might indicate a communication problem or poor pricing decisions, as utilization remains low in the face of availability or RTD. The calculations run automatically every day, week or month as appropriate. The results are then published in a browser-based HTML format, with the different reports (location, region, all) dynamically linked to allow quick comparisons as required. The appendices provide further Performance Monitor examples.

Status and Future Directions

Phase 1 as outlined is operational, providing objective results from multiple perspectives. The present design addresses seasonality through the ability to compare locations under similar market conditions. It also accommodates changes in market demand, as poor utilization and revenue performance could still indicate neutral performance on a relative standpoint. The location-specific characteristics allow the performance of a given location to be tracked (dynamically) in time, dynamically as the averages change over time.

Current efforts are focused on the opportunity costs and how they impact on the bottom line. Phase II, which is in development, quantifies the opportunity costs of the rate and ATD, in conjunction with some threshold level of utilization. The opportunity cost of rate and ATD are different depending on the level of utilization. If under-utilized the cost associated with turning business away may be far greater than in the case of decent utilization. Even with good utilization, however, there could still be a

cost associated with turning away high-paying 'last minute' consumers while renting to low-paying–early booking consumers. The opportunity cost is a function of the individual rates of the rentals, turndowns and utilization.

Table 2: Sample performance metrics

Utilisation Location	Fleet	On rent	Utilisation (%)	Scale
BWI	700	630	90	High
DEN	2000	1820	91	High
LAX	5000	4100	82	Low
EWR	2000	1650	83	Low
MCO	7000	5300	76	Low
PIT	500	415	83	Low
		Average	84	

Availability turndowns Location	Fleet	ATD	ATD/fleet (%)	
BWI	700	150	21	Low
DEN	2000	650	33	Low
LAX	5000	550	11	Low
EWR	2000	900	45	High
MCO	7000	1000	14	Low
PIT	500	415	83	High
		Average	35	

Rate turndowns Location	Fleet	RTD	RTD/fleet (%)	
BWI	700	100	14	Low
DEN	2000	1000	50	High
LAX	5000	2000	40	High
EWR	2000	750	38	High
MCO	7000	125	2	Low
PIT	500	125	25	Low
		Average	28	

Figure 2: Performance Monitor benchmarks
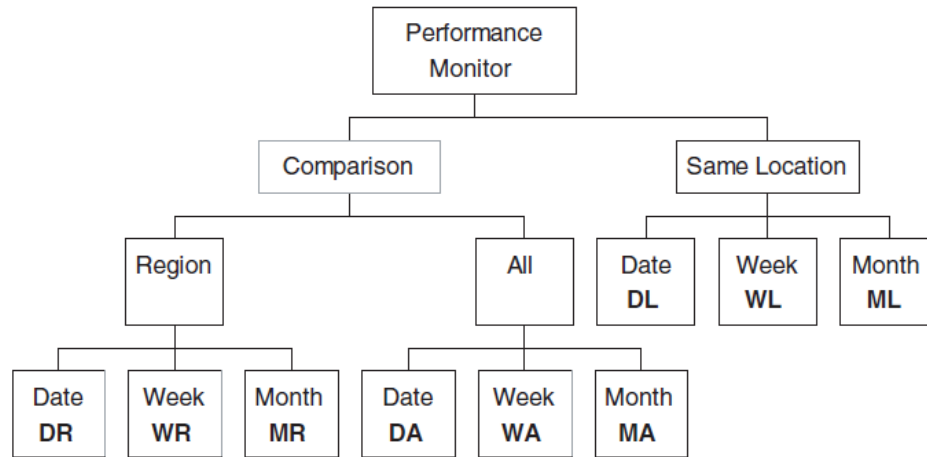
Table 3: Performance comparison

Location	Utilisation	ATD	RTD	Score	Comments
BWI	High	Low	Low	7	Could have yielded more, lost RPD
DEN	High	Low	High	8	Good
LAX	Low	Low	High	3	Should not have yielded as much, lost utilisation
EWR	Low	High	High	1	Blacked-out with idle cars
MCO	Low	Low	Low	4	Indifferent
PIT	Low	High	Low	2	Blacked-out with idle cars

Future developments beyond this second phase will look to include rate tracking and movements, and exogenous influences (customer service, market share) with current information on turndowns for improved elasticity estimates. The end goal is to understand the impacts of rate movements and controls upon performance to optimize fleet and pricing decisions further.

Figure 3: Performance Monitor output

Performance Monitor - WR

Weekly Ranking
for All Locations relative to the Regions' totals

Scale:	1 2 3	4 5 6	7 8
	Worse		Better

	Week							
Location	Week of 3/6/2000	Week of 3/13/2000	Week of 3/20/2000	Week of 3/27/2000	Week of 4/3/2000	Week of 4/10/2000	Week of 4/17/2000	Week of 4/24/2000
KOA	7	7	7	7	7	7	8	7
KOAN01	4	4	4	4	4	4	4	4
LIH	7	7	7	7	7	7	7	7
MKK	4	4						
OGG	4	7						
WAI	4	4						
BUR	4	2						
DEN	6	3						
FAT	2	2			6	6	1	1
LAS	4	7		7	7	7	3	7
LAX	4	4		4	4	4	7	7
OAK	4	4	4	2	2	4	2	2
ONT	7		2	2	5	2	2	4
PHX	1	1	1	6	1	3	6	6
PSP	8	8	8	4	7	7	5	2
RNO	6						4	4
SAN	6							4
SFD	4							7
SFO	5							7
SJC	4	4	5	5	2	5	4	4
SNA	7	7	5	5	4	4	2	2
TUS	7	5	7	4	4	4	3	7

This "Signals" that PHX has experienced several weeks of relatively low utilisation with High Rates and Blackouts

More recently, they still have High Turndowns, but the Utilisation has improved

References

Geraghty, M. K and Johnson, E. (1997) 'Revenue management saves national car rental', *Interfaces*, 27, 1, 107–127.

Cook, T. S. (1998) 'Sabre soars', *ORMS Today*, 26th–31st June.

Dargay, J. M. (1993) 'Demand elasticities: A comment', *Journal of Transport Economics and Policy*, 27, 1, 87–91.

Manugistics (2001) 'Pricing and Optimization at a Glance', http://www.manugistics.-com/documents/PRO_glance.pdf

Oum, T. H. (1989) 'Alternative demand models and their elasticity estimates', *Journal of Transport Economics and Policy,* 23, 2, 163–187.

PROS (2001) 'Energy Industry Showcased at PROS Pricing and Revenue Optimization Conference', http://www.prosrm.- com/news/2001conferenergy82201.html

Smith, B. C., Leimkuhler, J. F. and Darrow, R. M. (1992) 'Yield management at American Airlines',
Interfaces, 22, 1, 8–31.

Appendix 1

Appendix 2