

# Performance Monitor: The Opportunity Costs of Revenue Management

Chris K. Anderson, University of Western Ontario,  
Montgomery Blair, Dollar Thrifty Automotive Group

*Revenue management (RM) brings rational approaches to pricing for goods and/or services with a limited shelf life. The practice of dynamically pricing a perishable product across different market segments continues to be applied across an ever-increasing set of business arenas. While numerous consulting and software development firms preach the bottom line impacts of RM practice, little effort is applied to monitoring the success of RM systems once in place. The continued success of RM hinges upon the ability to link organizational performance to the pricing and capacity decisions of RM systems. This link both reinforces the financial gains attributable to RM and indicates opportunities for future improvement. This paper outlines Performance Monitor, a phased approach to performance measurement designed and implemented at Dollar Thrifty Automotive Group, Inc. The focus is on the impact of RM practice via a dissection of the lost revenue opportunities of historic decisions.*

## Introduction

Revenue management (RM) is being utilized in vast array of service industries; to allocate airline seats, rent cars and sell advertising space. Virtually any business setting that has a relatively high ratio of fixed to variable costs, products of limited shelf life and possesses the ability to monitor uncertain demand can profit from RM tactics. The processes of RM are attractive, since they provide rational approaches to dynamic pricing, enabling organizations to maximize revenues and increase profits. At its basic level, RM is simply market segmentation and inventory control. Market segmentation allows the determination of consumer's willingness to pay, with inventory control achieved through pricing changes across the different market segments.

There is considerable hype surrounding the success of RM. American Airlines attributes \$1bn in incremental revenue to its applications (Cook, 1998), while National Car Rental attributes its entire turn around to RM (Geraghty and Johnson, 1997). Those organizations that fail to practice RM while their competition does are at severe competitive disadvantage. Reflecting on the demise of his once successful airline, former PEOPLExpress CEO, Donald Burr 'What you don't know about Revenue Management could kill you!' (Bryan, 1989).

The original practice of RM began with overbooking, as airlines sold beyond capacity to alleviate low aggregate average yields (passenger miles travelled) resulting from customer cancellations and no-shows. Post-deregulation, as discounted sales became more popular, RM expanded to include market segmentation, demand forecasting and optimization. The end product of most RM systems was in the form of some type of 'recommendation'. Early on, they were in the form of availability of different fare/price classes, then dynamic pricing and now continue to evolve as the breadth of application evolves. Revenue management is growing into the inclusion of the supply side and even encompassing the larger 'enterprise' perspective. As the scope of RM increases, the need for a full understanding of the actions resulting from its implementation becomes paramount.

The practice of RM grows as success stories increase awareness, and new RM consulting firms are created and established firms offer services. The increased breadth of application requires more sophisticated science and tools. Revenue management service providers or system developers often outline the potential benefits of RM application through analysis of past performance. An analysis is usually performed to determine the amount of revenue and profit that can be attained by enhancing the systems. For a variety of reasons, especially the absence of a control, many of the same problems and poor assumptions that haunt assessment of RM potential, are present with the efforts of performance monitoring a posteriori (Blair and Anderson, 2002). The difference is that, in the initial stage prior to development, there is an absolute vested interest in the magnitude of the final result. If one can show that there lies a potential for expanding profits by over 100m, it nicely frames the cost of a system with one less digit. What is omitted is how a particular system will convert the potential profits. The assumptions are that the particular system will have generated the additional profit and that this additional potential profit (in the past) will be present in the future.

The RM systems that many organizations implement are just a part of a larger RM structure which includes the technical tools, limitations, decision support tools, human resource perspective, management structure and various logistic operational aspects of implementation. An individual system

alone may not result in the additional profit, but rather aid people with advanced tactics to capitalize on the potential.

Revenue management systems simultaneously use the limited, often censored, industry environment while influencing changes to it. A system that is developed for specific objectives may, through its use, alter the competitive environment. For example, many systems do not fully integrate or internalize competitive prices within their optimization, as it may not have been very feasible in the past. By using the systems' output over time, or as many companies utilize similar systems, however, the environment may evolve in which competition is much more (or less) important. In the car rental industry, erroneous assumptions on the price elasticity coupled with perceived risk-averse actions could have influenced overall industry rates down, greatly diminishing the profitability of all companies and the value of their systems' themselves.

In light of all the ambiguity, there is still inherent value in developing tools and enhancing RM systems. But more emphasis needs to be placed on the intricacies of quantifying the benefits gained from using these systems and the role each component plays. To make bold assumptions such as an 'improvement of X percent to the bottom line' is too vague and adds little substance to a meaningful analysis. Unless the amount of profit expected is fairly certain and all other influencing factors remain the same, this is merely a sales tool. Even if studies showed that implementing an RM system increased profitability, it should be important to determine which aspect of the system dominated that gain, and how much more could have been achieved. Taking this point one step further, even if the accountability could be traced to a particular system, questions should arise as to which mechanism had the strongest influence on performance, whether it can be replicated, and what happens if the environment changes. The underlying goal is a tracing back to the actions that created (or lost) value versus just a high level aggregation of the impacts.

### **A Focused Opportunity Cost System in the Car Rental Industry**

Although it may seem like an onerous task not only to quantify the benefits and determine the value, but also assess the link between RM systems and performance, there are methods that can shed light on this. One such method is under development at Dollar Thrifty Automotive Group, Inc. (DTAG), where the task of measuring performance is underway.

The performance monitoring system is being developed in multiple phases. The first phase (Blair and Anderson, 2002) is an objective-based scalar system which analyses relative performance from numerous perspectives. It can primarily be used to signal opportunities and trends from the overall

performance of locations. Phase II lies at the other end of the spectrum, as it will result in a dollar value of that which could have been done. It will supplement final accounting reports which show the revenue and profitability achieved by calculating the amounts that were possible. As a hypothetical illustration of this, assume Denver International Airport achieved revenue of \$10m with a profit of \$1m and a benchmark revenue per day (RPD) of \$35, but it is later calculated that it could have made \$12m with a profit of \$3m. By quantifying the potential, it will show the 'money that is left on the table'. It will also illustrate where the opportunity was lost and, thus, begin to shed focused light on what can be done in the future to capitalize on this. A major differentiating quality of this system can be appreciated through an extension of the previous hypothetical example. If Denver, in the same period the following year made only \$9m in revenue, \$500,000 in profit, and a similar or even lower RPD, but had little to no measurable potential above this, then RM pricing actions may have been good even in light of the perceived inferior performance. Exogenous forces such as competition or the macroeconomic environment could be to blame and not pricing. Likewise to this scenario, if that following year they achieved similar or even better numbers, but the percentage of potential was increased, it could be interpreted as a missed opportunity and poorer RM performance.

The simplest business model for some low-cost firms, such as DTAG, is to price extremely low and take as much business as possible by offering the best value. In fact, if prices are greater than costs, a firm would want to do this as much as possible and thus expand its profits with every sale. This strategy falters if the firm has a finite amount of product/service to deliver. If demand exceeds expectations or ability to handle daily supply fluctuations, the first reaction may be to increase rates on the remaining supply. If the firm still exhausts its supply, its price elevation has improved profitability. It still could have made even a greater amount if at any time prior to expiration; it turned away a consumer that was willing to pay an amount greater than that of its lowest priced sale. *The amount of this difference is the opportunity cost.* This concept is the heart of RM systems. Through forecasting, they attempt to anticipate not only the demand but also the time when it will materialize. The natural stochastic character of the demand ensures there will be a risk involved with turning away any of even the lowest priced demand. Even in light of anticipating ample demand at an elevated price, some might want to ensure some subsistence level of low priced business to diminish their risk.

Traditional RM systems, most notably airline RM systems, turn away low price demand or implicitly change prices by closing out availability of certain product (fare) classes. These availability controls represent discrete price changes and result from an industry's inability correctly to implement tactical reactions to a systematic approach of selling inventory, largely owing to the size of the problem

at hand and legacy information systems, although the advent of low-cost-carriers is migrating the airline pricing approach toward s that of car rental with more seamless continuous price changes.

During known busy travel dates, significant yielding still occurs. At a time when consumers are even less price elastic, it is often surprising how much of a discount is given to those who book their reservation early. This can be viewed in the car rental industry, where rates are lower three months prior to arrival, or in the airline industry, where deeply discounted fare classes have some level of allocated inventory. On 20th December, a consumer wishing either to pick up a car, to fly or to rent a room for the holiday period will undoubtedly experience 'high' rates. If the rates are too high, they may in fact forgo the purchase by finding a substitute or most likely end up paying the elevated amount. The majority of the bookings still come in the final weeks of the booking cycle, which to some may seem odd, as there is such an incentive to book early. This is not nearly as odd as the fact that more than likely someone who is flying on the same flight, picking up a car on the same day or staying in an adjacent room is doing so for far less. If a location runs out of cars and 'turns away' people likely to be willing to pay very high rates, while simultaneously renting to some at low rates, clearly there is an amount of lost revenue equal to the difference of the willingness to pay of the high valued consumer and the discounted rate. This could probably have been avoided with more information, superior systems or better human implementation. The goal is not to charge the last minute consumers more, but to reduce the dramatic discounts issued to the early business — as systems improve, the pricing response should become less reactionary and more proactive. Even for a company like DTAG, which strives to deliver value, they can actually achieve this through higher rates earlier which will result in more availability later when the competition is either 'blacked out' or gauging. The company can increase their perceived low-cost/low-priced relative placement through a more proactive approach. Increased value will come to a greater number of people while simultaneously generating more profit. From the basic laws of economics, this can be viewed as transferring the excessive consumer surplus of early booking customers into more quantity and better value to the later masses of consumers — a more efficient allocation of resources.

So why is this not really done? The exact level of price elevation is not known, misconceptions about supply (fleet) flexibilities, perhaps the misperception that low price equates to higher value when offered for a longer time and the alternative cost of spoilage is huge. This spoilage cost is not merely the difference of the rates, but it may be almost the entire amount of the rental itself. Whether for a 'known' high demand date such as 20th December or even in the instances of low demand dates, there could be opportunity costs associated with different pricing strategies for those who will travel.

Specifically with in the car rental industry, there are several challenging issues associated with analyzing and assessing R M performance. One of the main complexities associated with very dynamic pricing practices is that the n u m b e r of variables is overwhelming. For a single arrival date, there may be several car classes, different lengths of rent, distribution channels and corresponding rate codes. This complexity only increases with the practice of changing these rates almost daily along with the relative ranking and comparison with competitors' similar actions. To address this complexity, DTAG has approached performance measurement in a phased approach.

### **Performance Monitor**

Phase 1, detailed in Blair and Anderson (2002), is a relative comparison o v e r time and across locations and provides insight into how well (relatively speaking) a location is performing. Phase II, at the other extreme, directly quantifies the costs associated with past decisions for specific locations. As improvements are made and each of these advance with modified subsequent phases, a tracing back will occur and m o r e information as to the causality and relationships between actions and performance will become clear. A main tertiary benefit to this will undoubtedly be immediate enhancements to the RM system components themselves, especially to the business modelling and price elasticity knowledge.

Revenue management is a process of managing the trade-off between yield and spoilage. If discounting is too aggressive too early, yields are low; conversely, if discounting is curtailed, sales may be low and inventory wasted. The practice of RM uses a combination of historic demand and current information to allocate discounts (or set the amount of the low -priced market segment business) to utilize the inventory best, while maximizing revenues. Revenue management then uses overselling to help mitigate mistakes. Overselling (Biyalogorsky *et al.*, 1999) allows sales to higher-yielding customers in the anticipation that some booked customers may not materialize or that some low -yield customers may be persuaded, via compensation, to delay their reception of the good and/or service. Performance Monitor looks to measure this trade-off, specifically looking to quantify the financial impacts of spoilage and a lack of yielding (discount limitation).

Attempts to understand spoilage or spill, as it is often termed, are not new. T h e airlines assess the impact of lost demand in a two -step process. Initially, reservation histories are unconstrained to determine a measure of demand, typically done under the assumption of a Gaussian distribution (Belobaba, 1987; Lee, 1990). Unconstrained demand estimates  $f(x)$  are then used to calculate the expected demand beyond capacity  $C$  as an estimate of expected spill  $S P$  (Swan, 1979):

$$SP = \int_{x-C}^{\infty} f(x)(x - C)dx$$

Once spill has been estimated, a cost is associated with the spill by estimating a spill fare and multiplying the spilled demand by this fare. Swan (1992) indicated a weighted average fare (0.8 to lowest fare and 0.2 to overall average fare) as a good approximation to the spill fare. Belobaba and Farkas (1999) indicated that, as a general rule, the spill fare should be below the flight's average fare. These spill estimates focus on turned away demand and apply an aggregate fare to all this demand. An attempt is made partially to disaggregate spill, looking at how utilization impacts this estimate.

### **Performance Monitor Phase I**

The inputs for Phase I are utilization and turndowns. Utilization, the fraction of available fleet on rent, provides an indication of spoilage. Turndowns, both availability (lack of capacity) and rate (available capacity, but not the right price) start to provide insight into yielding. Rate turndowns (RTD) reflect the business interested but not willing to pay the current price. Availability turn downs (ATD) reflect an apparent lack of capacity.

Phase I is a comparative tool. Instead of qualifying exactly what levels of utilization and turndowns are strong/poor; relative measures are used, resulting in performance on utilization and turndowns judged among DTAG 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.

Top -tier locations do an excellent job of achieving high utilization while not turning away high value traffic. 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 customers.

**Table 1: Performance quality grid**

<i>Utilisation</i>		<i>ATD</i>	<i>RTD</i>	<i>Comments</i>	<i>Score</i>
High	Top	Low	High	Yielded nicely	8
High		Low	Low	Could have yielded more	7
High		High	High	Could have definitely yielded	6
High	Mid-tier	High	Low	Turned away high value	5
Low		Low	Low	Indifferent — Too much fleet?	4
Low		Low	High	Yielded — cost rentals	3
Low	Bottom	High	Low	Turned away business	2
Low		High	High	Very poor	1

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 u n d e r 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.

### **Performance Monitor Phase II**

The inputs for Phase II are the same as Phase I: utilization and turndowns. There are two outputs from this phase: the quantity of spoilage and an independent opportunity dollar cost amount. The opportunity dollar value will be an approximation of the additional revenue that was available and likely to materialize if better RM pricing actions had been implemented. This model may err on the conservative side and should be used as a floor for the true potential. Likewise, although the results are actual revenue and profit values, additional insight may be gamed from analyzing the trends and ratios of the results.

Logically, for every turndown, based on rate or availability, there may be an opportunity cost of that neglected service. For example, if a consumer is willing to pay \$40, but no car is available, this is referred to as a ‘black - out’ and recorded as an availability turndown. The difference between this rate of \$40 and the lowest rate that another consumer who probably booked months in advance, say \$25, is defined as an opportunity cost and would be equal to the difference of \$15. If there were cars that sat idle then the opportunity cost would be the entire amount of \$40.

Similarly, if a consumer shopped days prior to arrival and was quoted a rate of \$40, which they declined to accept, there may be a similar opportunity cost, somewhere between \$0 and \$15 (\$25 car) for the days with decent utilization and up to just under \$40 (as high as the consumers willingness to pay) if there was spoilage. This, by definition, is a ‘rate turndown’ and, although the consumer was n o t



willing to pay this amount, there is a good chance they were willing to pay more than the lower-priced \$25 renter.

To scale up this analysis, the technical challenge is not so much centered on the rate turndown assumptions, as on the lower rates actually taken. For example, in the above illustration, the substitution of the lower rental for the higher one is preferred, but what about the next turndown? Ideally, a matching of the turndowns to individual rentals would be most appropriate, but the approximation suggested is to use the mean and standard deviation of the rentals as a proxy for actual rates.

By calculating the mean and standard deviation of rentals, a good approximation for the rates taken can be assumed. For example, if the mean rates for a particular arrival date, car class and length of rent is \$30 with a standard deviation of \$8, then it can be assumed that there were probably several rentals taken at rates around \$22. This should intrinsically take into account the actual rate as well as the weight or amount of business. If, for example, a location offered just a few rates at \$20 then quickly raised it to \$30, the statistics would reflect this. Taking this one step further, the proxy for low rates will be the mean less one standard deviation. The associated opportunity cost will be dependent on whether or not there were idle cars and will follow the schedule depicted in Table 2.

The amount of the opportunity cost depends significantly on whether there were idle cars, referred to as spoilage. Achievement of 100 per cent utilization may be unattainable, perhaps only coming at the expense of customer satisfaction or logistical improvements. Therefore, some threshold utilization is installed, anything less than this represents idle inventory, above being deemed acceptable. Consider a threshold utilization of 85 per cent, selected somewhat arbitrarily where anything less than this is simply not good. If there was not enough demand, then this quantity may be due to over-fleeting. Over-fleeting on a very granular level may not be avoided as to ensure ample fleet for higher demand periods within the cycle time to move or acquire inventory. Nonetheless, if it is present and any demand is turned away, there is the potential for a high opportunity cost. This opportunity cost may be broken down to either a 'fleet' or 'utilization'. This dissection is dependent on the demand forecast and will help lead to some accountability. Table 3 outlines the determination of the number of vehicles with which an opportunity cost may be associated.

The three hypothetical examples outlined in Table 4 illustrate opportunity quantities with an available fleet of 1,000 cars, with different fleet quantities and utilization quantities, yet the same total opportunity quantity (OQ).

**Table 2: Opportunity cost approximations**

Availability turndown = ATD	Rate turndown = RTD
No. of ATD = $B$	No. of RTD = $S$
No. of rentals = $N$	Rate = $R_i$
Mean of all rates = $X$	Standard deviation = $\sigma$
Low utilisation $_l$	High utilisation $_h$
$ATD_l Cost_l = X + \sigma$	$ATD_h Cost_h = 2\sigma$
$RTD_l Cost_l = \min[R_i, X] - \sigma$	$RTD_h Cost_h = \max[0, R_i - (X - \sigma)]$

**Table 3: Opportunity quantity**

UCFD — unconstrained forecasted demand	
Fleet — available vehicles for rent by car class	
Fleet quantity (FQ)	$\max[0, 0.85 \times \text{Fleet} - \max(\text{Rentals}, \text{UCFD})]$
Utilisation quantity (UQ)	$\max[0, \min(0.85 \times \text{Fleet}, \text{UCFD}) - \text{Rentals}]$
Opportunity quantity (OQ)	$\max[0, (\text{FQ} + \text{UQ})]$

The separation of fleet and utilization opportunities allows the allocation of costs to different parties. Given the dynamic pricing nature of the car rental industry, it is often inferred that low yields are attributable to pricing errors. A joint consideration of OQ with turndown information provides an insight into lost revenue opportunities. Table 5 illustrates two hypothetical examples separating turndown information in accordance with fleet utilization. Turndowns, both rate (RTD) and availability (ATD), are separated into low Q and high (h) as measured relative to utilization. Costs associated with each of these turn downs, as depicted in Table 2, differ depending on whether the turndown is associated with low utilization (spillage) or not.

The next step is to add in probabilities or scaling to each component. This will help to adjust for turndown sampling issues and enable the user to adjust the weight in accordance with the nature of turndowns. In other words, if there is further investigation into RTD, and it is determined that some of these RTD are consumers that ultimately returned and booked, or diverted to a different rate class, we may want to diminish the weight of this opportunity cost. Other issues may be the limited distribution source information on turndowns currently received. Reservations arrive through numerous distribution sources, whether they are from a GDS, 800 numbers, company specific internet distribution or third-party internet facilities. If turndowns only come from one source such as the phone system, perhaps they should be treated as a proxy for a larger quantity.

**Table 4: Fleet, utilisation and opportunity quantity**

<i>Example 1</i>			
	Fleet	1,000	
$FQ = 850 - 750 = 100$	Threshold utilisation	850	
	Unconstrained demand	750	$OQ = 150$
$UQ = 750 - 700 = 50$	On rent	700	
<i>Example 2</i>			
	Fleet	1,000	
$FQ = 850 - 700 = 150$	Threshold utilisation	850	
	Unconstrained demand	675	$OQ = 150$
$UQ = \text{Max}[0, -25] = 0$	On rent	700	
<i>Example 3</i>			
	Fleet	1,000	
$FQ = \text{Max}[0, -50] = 0$	Threshold utilisation	850	
	Unconstrained demand	900	$OQ = 150$
$UQ = \text{Max}[0, 150] = 150$	On rent	700	

**Table 5: Turndown separation, low versus high utilisation**

ATD <sub>l</sub> = Min[OC, B]		B total number	ATD
ATD <sub>h</sub> = Max[0, (B – ATD <sub>l</sub> )]		S total number	RTD
RTD <sub>l</sub> = Min[S, Max(0, OC – ATD <sub>l</sub> )]			
RTD <sub>h</sub> = Max[0, S – RTD <sub>l</sub> ]			
<i>Example 1</i>			
OQ = 100	ATD <sub>l</sub> = Min[100, 10] = 10		
B = 10	ATD <sub>h</sub> = Max[0, (10 – 10)] = 0		
S = 50	RTD <sub>l</sub> = Min[50, Max(0, 100 – 10)] = Min[50, 90] = 50		
	RTD <sub>h</sub> = Max[0, (50 – 50)] = 0		
<i>Example 2</i>			
OQ = 10	ATD <sub>l</sub> = Min[10, 50] = 10		
B = 50	ATD <sub>h</sub> = Max[0, (50 – 10)] = 40		
S = 100	RTD <sub>l</sub> = Min[100, Max(0, 10 – 10)] = 0		
	RTD <sub>h</sub> = Max[0, (100 – 0)] = 100		

The probability that a turndown may have materialized will depend on the type of turndown and utilization of the day/class/length-of-rent in question. For example, if cars are idle (low utilization), there is not only a higher cost associated with the turndown, but also a stronger likelihood that the revenue could be captured. Table 6 combines the scaling with the number of turndowns. The following details an example of opportunity costs associated with rental agreements for cars at Denver International Airport.

### Performance Monitor Example

In 2001, 22nd March was a strong rental day for DTAG, with approximately 94 per cent of the available fleet at Denver on rent. As detailed in Table 7, revenues from C class cars (compact) were \$4,327. Given the presence of turndowns, it is argued that more revenue could have been generated

through renting the same number of cars at higher rates. Early rate reductions created potential opportunity costs of 13 per cent incremental revenues equivalent to a two fold increase in profits.

The location offered some low rentals, then yielded resulting in a mix of rates with an overall average of \$25.95. As the yielding occurred, the location received some RTD. Ultimately, there were no longer any cars available, and the location probably 'blacked out'. This resulted in last minute enquiries or ATD and a significant opportunity cost. Ideally, some of the early discounts should have been avoided, enabling rentals at higher rates closer to the arrival date. Since utilization was high, exceeding the 85 per cent threshold, opportunity costs were lower, associated only with yielding and not spoilage.

#### *Availability Turndowns*

Eight requests for rentals were not fulfilled. By not offering a rental to these last minute customers, DTAG gave up revenue equal to the difference between that requested and that received. These ATD occurred at a similar time to the rentals filled at \$47.99 rates. For the first one of these, the (daily) opportunity cost could have been as much as the difference between this \$47.99 and the lowest rental of \$14.66. This would have been a rate difference of \$33.33 across a length-of-rental (LOR) of four days, for \$133 in additional profit. The technical challenge would be to then take the next availability turndown and perform a similar calculation with the next to lowest rental and so forth. Although this may be the ideal measurement of this opportunity cost, it poses technical problems. The model conservatively estimates these costs. By using a value of  $2\sigma$ , one can represent a conservative approximation to this difference for the entire count of ATD. Not only will this proxy contain the high and low rate values achieved, but it will also weight the number of occurrences. For this example, each availability turndown would translate into \$16.08 of forgone revenue, ie

$$ATDC = ATD_h = 2\sigma = \$16.08$$

With two ATD occurring with LOR of four days, and six with LOR of three days, the total forgone revenue from ATD exceeds \$418 [ $16.08(4 \text{ days} \times 2 + 3 \text{ days} \times 6)$ ].

**Table 6: Opportunity cost scaling**

	No.	Cost	Scale/probability	Total cost
ATD	ATD <sub>l</sub>	$X + \sigma$	$\lambda_1$	$\lambda_1 \times (X + \sigma)$
	ATD <sub>h</sub>	$2\sigma$	$\lambda_2$	$\lambda_1 \times 2\sigma$
RTD	RTD <sub>l</sub>	$\text{Min}[R, X] - \sigma$	$\lambda_3$	$\lambda_2 \times \text{Min}[R, X] - \sigma$
	RTD <sub>h</sub>	$\text{Max}[0, R - (X - \sigma)]$	$\lambda_4$	$\lambda_3 \times \text{Max}[0, R - (X - \sigma)]$

**Table 7: Compact class rentals — 22nd March, 2001, Denver International Airport**

Available fleet 503 on rent 473	Rental count 48 rentals
LOR category L (3–4 days)	Rate range (\$/day) \$14.66–47.99
ATD 8 (6 LOR 4 days, 6 LOR 3 days)	Total revenue \$4,327
RTD 4	
3 at \$21.99/day	Mean rate (X) \$25.95
1 at \$47.99/day	Standard deviation ( $\sigma$ ) \$8.04

### *Rate Turndowns*

Rate turndowns occur at a known rate, with an unknown willingness to pay of the potential customer. In this example, four RTD were received, three were at \$21.99/ day and one at \$47.99. The forgone revenue, given strong utilization, is the difference between what the customer would have paid (unknown) and low priced rentals. Similarly to the ATD, an attempt is made to model this cost through the difference between the individual rate of the turndown and a weighted approximation of low rates. The weighted approximation of low rates is the mean less one standard deviation or, in this case, \$25.95 — \$8.04 or \$17.91, resulting in an opportunity cost of \$4.08 for the \$21.99 turndowns and 830.08 for the \$47.99 rate.

$$\begin{aligned}
 \text{RTD} \times C &= \text{RTD}_i = \text{Max}[0, R - (X - \sigma)] \\
 \text{Max}[0, \$21.99 - (\$25.95 - \$8.04)] &= \$4.08 \\
 \text{Max}[0, \$47.99 - (\$25.95 - \$8.04)] &= \$30.08
 \end{aligned}$$

The resulting opportunity cost from the four RTD is \$139.20 (three at \$4.08/day for four days, one at \$30.08/day for three days), producing a total opportunity costs, both RTD and ATD, of \$557.28.

This means that the location could have achieved an additional \$557.28 in revenue for the same costs and rentals, ie by offering low ‘discounted’ rates in advance, it cost the company \$557.28 of forgone revenue.

It may be important to view this from two different perspectives. First, for this specific ‘piece’ of business there were 48 rentals totaling \$4,327. The common benchmarks would show that better pricing would have resulted in an average yield gain of 13 per cent, or assuming a profit margin of 12 per cent, would have increased profits by over 100 percent. The other perspective worth mentioning is that of scale. Although this \$557.28 may not seem like much, it is only one car class, one length of rent category, for one day, in one location.

Looking at a few length of rent categories, D (1–2 days), L (3–4 days), W (weekly) and X (8 days or more) as well as the three major car classes C1 (compact), I (intermediate) and M (midsize) for the

week of 16th to 22nd June, 2003, at Denver, one can get an idea of the scale of revenues. Table 8 summarizes utilization for the week, indicating that about 90 percent of the available fleet was on rent. Table 9 indicates total opportunity costs as a percentage of total sales, and Table 10 indicates that the bulk of these opportunity costs were associated with ATD, which is expected given the high utilization. Total revenue for this associated week and these three car classes at Denver was \$172,000 resulting in a total lost opportunity of almost \$12,000. This should indicate that, while DTAG had a great week at Denver, there are still significant revenues to be captured.

**Table 8: Utilisation rates, Denver International Airport**

	<i>16th June</i>	<i>17th June</i>	<i>18th June</i>	<i>19th June</i>	<i>20th June</i>	<i>21st June</i>	<i>22nd June</i>	<i>Week</i>
Utilisation	0.90	0.90	0.92	0.89	0.91	0.86	0.91	0.90

**Table 9: Total opportunity costs as percentage of total revenue**

<i>Car class</i>	<i>LOR</i>	<i>16th June</i>	<i>17th June</i>	<i>18th June</i>	<i>19th June</i>	<i>20th June</i>	<i>21st June</i>	<i>22nd June</i>	<i>Week</i>
C	D	4.1	0.0	8.1	61.0	65.5	18.7	30.9	27.1
	L	2.2	5.1	11.9	11.7	38.6	2.1	7.5	7.8
	W	12.7	4.2	23.8	11.3	15.6	0.0	39.2	9.9
	X	0.0	4.8	0.2	6.4	0.7	2.7	5.0	2.8
I	D	12.4	16.3	67.2	47.0	98.8	25.9	29.7	32.0
	L	5.0	5.8	5.7	20.1	10.9	3.2	8.8	8.3
	W	5.3	9.4	4.4	2.6	2.3	1.5	8.5	4.3
	X	0.1	3.4	3.4	6.5	3.2	0.3	9.8	3.6
M	D	0.0	0.0	0.0	86.4	0.0	0.0	0.0	11.9
	L	0.0	3.2	2.6	9.2	0.0	6.5	5.8	4.2
	W	5.1	6.3	18.9	12.2	2.5	5.4	3.3	7.0
	X	4.8	21.3	1.9	2.6	6.4	3.5	7.4	4.3
Total		4.5	7.0	6.6	10.6	6.1	3.6	8.9	7.0

**Table 10: Total opportunity costs from ATD as percentage of total revenue**

Car class	LOR	16th June	17th June	18th June	19th June	20th June	21st June	22nd June	Week
C	D	0.0	0.0	0.0	61.0	65.5	18.7	20.3	23.4
	L	0.0	5.1	11.9	11.7	33.2	2.1	0.0	6.0
	W	12.0	4.2	19.7	11.3	15.6	0.0	24.8	8.5
	X	0.0	4.8	0.0	6.4	0.7	2.7	5.0	2.8
I	D	8.6	12.0	25.0	47.0	70.9	25.9	0.0	17.7
	L	1.8	5.5	2.6	20.1	10.9	3.2	1.7	6.2
	W	4.4	8.9	4.4	2.6	2.1	1.5	3.4	3.3
	X	0.1	3.4	2.7	6.5	2.6	0.3	4.2	2.9
M	D	0.0	0.0	0.0	86.4	0.0	0.0	0.0	11.9
	L	0.0	3.2	1.7	9.2	0.0	6.5	5.8	4.0
	W	5.1	6.3	17.5	12.2	2.5	5.4	3.3	6.8
	X	4.8	21.3	1.9	2.6	6.4	3.5	7.0	4.3
Total		3.5	6.6	4.6	10.6	5.5	3.6	3.5	5.7

*Figure 1: Performance Monitor Phase II summary*

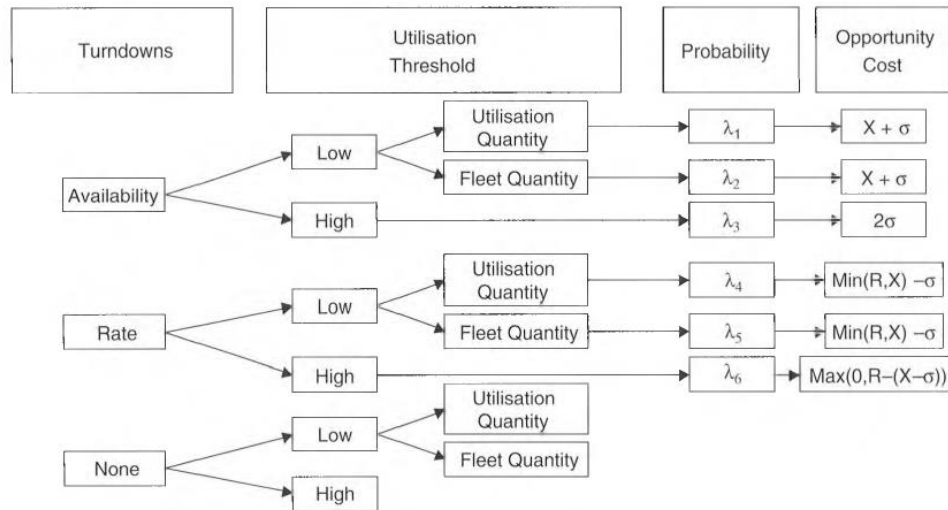


Figure 1 illustrates the relationships of the different segments of Phase II. This example illustrates a point in time when the location achieved decent utilization and was thus most likely well managed, with performance assessment from traditional methods (utilization and RPD) indicating strong performance. The fact is it could have done better, and this attempts to quantify the dollar amount. In the example, the opportunity costs were not scaled, as displayed in Figure 1, to reflect their likelihood of materializing or data limitations (not all turndowns recorded). The scaling to represent all business depends upon the degree to which turndown information is captured (relative to all business). At present, probability estimates for the different types of opportunity costs are being improved. Moving

forward, for the same day of week, day of year, one may see this performance and choose not to offer rates below \$20.99, or realize the size of likely turndowns and raise prices earlier as signaled by current turndown information. If the amount of actual revenue to that of the potential or opportunity costs increases, it could be assessed that performance was better, even if the RPD or other benchmarks are not indicating the same thing. Keep in mind that the risk of not offering the low rates to the early booking market segment is that it will result in an idle car, which would be quantified as a large cost and represented accordingly in the model.

## **Conclusions**

Performance Monitor is a first attempt to distill the impacts of RM actions upon performance. Novel to Performance Monitor is the use of lost sales information. Lost sales or turndowns provided a wealth of insight into historic pricing decisions if systematically captured and accurately recorded. Opportunity costs originate with turndowns, rate and/or availability. These turndowns are segmented into those that occur during periods of low and high utilization. Generally speaking, turndowns in periods of low utilization are more costly in terms of potential revenue and have a higher probability of materializing; those that occur in periods of improved utilization are less likely to be captured and potentially less costly.

After an arrival date passes, with perfect information one could pick and choose the reservation requests that would have resulted in the maximum revenue. If known in advance, the decision would not be that of not offering the lowest paying consumers availability, but offering it to them at the elevated price. If spoilage had occurred, turning away anyone away could have been potentially very costly and turning away consumers in excess of lower rates (defined here as mean minus one sigma) would be ill-advised.

The next generations of RM systems will incorporate forward - looking tools with similar ingredients of analyzing the business not taken or 'turned away' as well as that which materializes. This should result in lower opportunity costs and more profit. Understanding and quantifying the results of pricing actions will more appropriately frame the likely benefit directly relating to specific actions and systems. Poor secondary metrics and benchmarks will be either supported or disputed, given this performance metric. Highlighting exogenous factors, ascertaining acceptable degrees of risk and overall potential value from expensive RM system development or purchase can become less abstract.

The importance of tracking accurate lost sales data cannot be overstated. Lost sales or turndown data help to paint the 'what if' picture which may otherwise be lost. Moving forward, an



increasing percentage of reservations are completed online, though company specific websites such as [www.Dollar.com](http://www.Dollar.com) or third-party sites such as Travelocity and Expedia — it is equally important to monitor conversion ratios on these sites, ie the fraction of traffic that books a reservation, as these data become analogous to turndowns.

As this monitoring system advances, the systems, including the people using them, will be motivated to decrease the opportunity costs, which should become synonymous with better forecasting and early pricing decisions. Performance Monitor, while heuristic in nature, provides vital performance metrics essential for the continued success of pricing decisions at DTAG. Performance Monitor signals missed revenue opportunities and, more importantly, provides a measure of potential revenues that may be captured with improved RM approaches. The utilization of opportunity costs in conjunction with the relative benchmarking of Performance Monitor provides a solid foundation for the assessment of RM performance at DTAG or at any other firm practicing RM.

## References

- Belobaba, P. P. (1987) 'Airline yield management: An overview of seat inventory control', *Transportation Science*, 21, 1, 63—73.
- Belobaba, P. P. and Farkas, A. (1999) 'Yield management impacts on airline spill estimation', *Transportation Science*, 33, 2, 217-232.
- Biyalogorsky, H., Carmon, Z., Fruchter, G. E. and Gerstner, E. (1999) 'Research note: Overselling with opportunistic cancellations', *Marketing Science*, 18, 4, 605- 610.
- Blair, M. and Anderson, G. K. (2002) 'Performance Monitor', *Journal of Revenue and Pricing Management*, 1, 1, 57—66.
- Bryan, J. A. (1989) 'Donald Burr may be ready to take to the skies again', *Business Week*, 16th January.
- Cook, T. M. (1998) 'SABRE soars', *OR/MS Today*, June, 26-31.
- Geraghty, M. K. and Johnson, E. (1997) 'Revenue management saves national car rental', *Interfaces*, 27, 1, 107-127.
- Lee, A. O. (1990) 'Airline reservations forecasting: Probabilistic anti statistical models of the booking process', PhD dissertation, Flight Transportation Laboratory Report R90-5, Massachusetts Institute of Technology, Cambridge, MA.
- Swan, W. M. (1979) 'A systems analysis of scheduled air transportation networks', PhD dissertation, Flight Transportation Laboratory Report R79-5, Massachusetts Institute of Technology, Cambridge, MA.

Swan, W. M. (1992) 'The value of yield management information for routing and scheduling',  
Presentation to AGIFORS Yield Management Study Group, Brussels, 4 - 6<sup>th</sup> May.