10

Measuring Airline Profitability

Ben Baldanza
Chief Operating Officer
Grupo Taca
San Salvador
El Salvador

Understanding the profitability of an airline network is a complicated and often misunderstood process. Although total revenues and costs are generally adequately collected, defining which costs and which revenues to ascribe to subsets of the whole (hub, city, route, aircraft, and so forth) is necessary to truly understand where the company is making and losing money. Having this information is essential to make the important decisions related to fleet assignment, capacity planning, and route selection.

The Fundamental Problem in Measuring Airline Profitability

An airline's costs are driven principally by the its fleet selection and flight schedule. Fleet selection determines the majority of the capital costs as well as drives fuel burn rates and many maintenance items. The operating costs of the airline are a direct function of the schedule structure. The ratio of departures to block time, stage length, and the basic structure of flights (point-to-point or hub-and-spoke) determine the rest of the maintenance drivers, labor staffing levels, and non-aircraft capital costs. Although general overhead and varying productivity levels contribute to the overall costs, the choice of fleet and schedule structure form the basis for the airline's cost structure.

Airline revenues, on the other hand, are completely unrelated to this structure for any individual airline. Airline prices are often market competitive or regulated by a formula based on mileage. Prices set in either manner are based on the origin & destination (O&D) of the customer. Consider the following simple example:

Passenger origin & destination (O&D)	Syracuse, New York (SYR)–Los Angeles, California (LAX)
Nonstop mileage	2,345
Competitive fare	$250 one way (OW)

Three airlines offer service on this route, as shown in Table 10-1. Obviously, the cost of providing service between Syracuse and Los Angeles is different for each of these three airlines. Yet, owing to the pricing environment, the revenue earned for servicing a passenger travelling on Airline A, B, or C is identical, and it is based on factors unrelated to the given airline's fleet choice and schedule structure. This relationship is common for virtually every airline in the world except for the few purely monopolistic, price-protected markets that are still left. Thus we can sum up the fundamental problem for airline profitability measurement as follows: *Airline costs are based on flight segment as a function of fleet and schedule, while passenger O&D drives the airline revenue.* As a result, airlines must find a way to measure costs effectively for their fleet and schedule and to relate costs to revenues mapped over this cost structure. These tasks constitute the measurement of airline profitability, and they are the focus of the rest of this chapter.

Airline Profitability versus Airline Accounting

Profit equals revenue minus cost. Airline accounting departments collect cost and revenue data to develop formal financial statements. These data, although crucial to the development of meaningful profitability information, are not complete for measuring profitability, and if they are used for such they will likely lead to poor decision making. Exactly which revenues and which costs to consider when studying the profitability of subsets of the airline network is not intuitive, and typical accounting statements do not produce relevant data.

To study profitability, one needs to build data that support *decision making*, and this differs from the data used for *financial reporting*. The distinction is shown in Table 10-2.

Although formal financial reporting is necessary to

Table 10-1 Details Differentiating Three Airlines Serving the Syracuse, New York/Los Angeles, California Market

Airline	*Service Offering*	*Mileage*	*Aircraft Used*
A	Nonstop	2,345	Leased A-320
B	Jet-to-jet connection at a mid-continent megahub	2,900	Owned and depreciated 737 to hub, owned 767 from hub to LAX
C	Turboprop connect at closer regional hub	2,500	Contract ATR-42 from regional service provider to hub, leased 727 from hub to LAX

Table 10-2 Distinction in Data Collection to Support Financial Reporting and Decision Making

Area	*Financial Reporting Data*	*Decision Making Data*
Revenues	Accurate collection and reporting of realized revenues for a specified time period	Incremental revenue associated with a specific piece of the airline network, that is, *what revenues are at risk if the target network subset is changed*
Costs	Accurate collection and reporting of operating and non-operating expenses, proper accounting treatment of out-of-period expenses, accrual for long-term events (such as heavy maintenance)	Incremental costs associated with a specific piece of the airline network, that is, *what costs are actually affected if the target network subset is changed*

understand and correctly interpret the firm's financial position, airline profitability data must be compiled in such a way that it is possible to interpret the profit contribution of each piece of the network correctly. Typical accounting rules, and by extension accounting departments, are not geared to answer such cost and revenue questions. Therefore an airline that wishes to truly understand the profitability of each network component must develop a separate but related airline profitability measurement process.

Levels of Profitability Measurement

As we have said, it is essential to measure the profitability of each piece of the network. With that idea in mind, subsets of the total airline profit and loss (P&L) statement could include

- a single hub operation
- a specific city on the network
- a single nonstop market
- a specific flight within a market

Other breakdowns of network profitability could also be done. For example, items such as profitability by customer or by distribution channel also are of use to the airline. Although these are valid and complex issues in their own right, we will focus here on profitability by subsets of the flying network.

This level of profitability analysis is meaningful for any airline in any situation, and it can be created with basic accounting data and some creative effort from an airline's analytical and information technology (IT) staffs. We can take this as a "starting point" for a meaningful understanding of system profitability, and later we can enhance the analysis with more uniquely appropriate measurements.

An Overview of Flight-Level Profitability

One building block of any profitability analysis is the measurement of an individual flight. After a detailed review of the flight-level measurement process, we will discuss the appropriate aggregation of these data to provide market, city, and hub information.

To begin, we will define a simple airline structure. Consider the schedule shown in Figure 10-1. We will focus on the measurement of a single one-way SYR–OMA flight, carrying only passengers. Start by measuring incremental revenues and incremental costs for this flight.

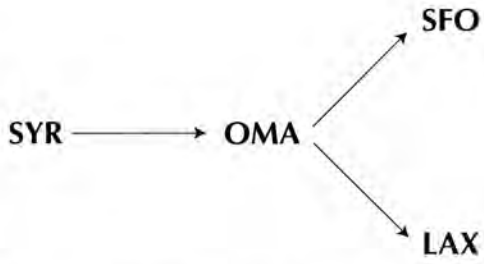

Figure 10-1 A Sample Schedule.

Incremental Revenues for SYR–OMA

There are two sources of incremental passenger revenues associated with the SYR–OMA segment: revenue from passengers boarding in SYR and deplaning in OMA ("local revenue"), and revenue from passengers boarding in SYR and connecting in OMA to either LAX or SFO.

The first source of revenue is straightforward, and these data are probably already collected by the accounting department. The second source of revenue is more complex. A proration methodology is needed to spilt the O&D fare into its two component parts, and the reason for this will be seen shortly. For now consider that any meaningful proration is acceptable (mileage, ratio of full fares, and so forth).

These are the basic components of revenue needed for undertaking the profitability analysis, although a simplifying assumption is implicit: 100 percent of the revenues we have defined are considered incremental. This means, for example, that if the particular SYR–OMA flight being examined were cancelled, this analysis would assume that the airline would lose 100 percent of the local revenue and 100 percent of the prorated portion of connecting itineraries that use the SYR–OMA segment. But the airline may not lose its local revenue and prorated revenue if, for example, the airline has multiple flights in the SYR–OMA market and could therefore possibly recapture some of the revenue lost from cancellation of this individual trip. For now, however, we will ignore this possibility and assume that we have the true incremental revenues.

Incremental Costs for SYR–OMA

The cost development for SYR–OMA has many levels. There are variable operating costs, aircraft ownership costs, overhead and non-operational costs, and network costs. To begin, we consider the variable operating costs of flying this one trip:

- fuel cost;
- crew costs;
- direct maintenance cost;
- servicing costs (ground handling, airport processing, landing fees, and so forth); and
- passenger variable costs (food, insurance, commissions, and so forth).

Together, this group of costs is referred to as the *variable costs* for the trip.

Second, we should consider *aircraft ownership* costs. In the very short term, ownership costs are not incremental and therefore are not relevant for some analyses (for example, cancelling one SYR–OMA trip in a year that does not reduce the airline's capital commitment for the aircraft used). Aircraft ownership costs will be discussed in detail below; for now assume that we can ascribe an hourly lease rate to the aircraft and "charge" SYR–OMA for the time it uses this equipment. The airline's treasury department will be useful for carrying out this analysis, since it will have the details on the aircraft ownership or lease contracts.

Third, we can allocate a portion of the airline's overhead and non-operational costs to this flight segment. As with aircraft ownership, the relevance of these costs is a function of what period of time and what aggregation is being examined. Assume for now that these costs are allocated back to the SYR–OMA segment on an available seat miles (ASMs) share basis. The accounting department can do this calculation easily.

Fourth, and finally, we must consider the costs of this trip to the rest of the network. Here we come to both real and opportunity costs. The real network costs of SYR–OMA include the cost of processing SYR connecting passengers at OMA, as well as the incremental cost of having those passengers on what otherwise would be empty seats on the connecting segment. The network opportunity cost is the most theoretical component of the flight profitability measure. It attempts to model the revenue lost by connecting passengers in OMA from SYR in seats that could have otherwise been sold to local OMA passengers or alternative connection points (in our simple system, only a passenger choosing to connect in OMA to fly between LAX or SFO). This will also be discussed in detail later in the chapter.

In summary, the incremental cost for SYR–OMA can be defined at several levels:

- variable flight costs only;
- variable flight costs plus aircraft ownership;
- fully allocated flight costs (variable plus aircraft ownership plus allocated overhead and non-operating costs); and
- any of these three levels with or without the contributing network-related costs.

Putting the Revenues and Costs Together

By combining the formats of revenue and cost, we can derive a series of measurements for SYR–OMA, each of which has a use based on a specific set of circumstances. (See Tables 10-3 and 10-4.) The choice of the appropriate cost basis is a function of what decision needs to be made and the time frame of the analysis. In general, a fully allocated measurement is most relevant with a longer lead time and/or greater aggregations of flights.

Figuring Network Contribution

To each of the three basic measures—variable segment profitability, variable profitability plus ownership, and fully allocated segment profitability—one of two network connectivity components can be added, as shown in Table 10-4.

Unlike with segment costs, choosing the network connectivity costs to ascribe to a flight segment is not a function of the business decision but rather a function of the availability of appropriate measures. The more robust network cost measurement (including opportunity costs) is always preferable if it is accurate, but many airlines do not have the ability to accurately measure this term.

This creates a range of measurements that are useful in different settings. Since the terms tend to get long quickly ("Variable profitability with ownership plus network contribution with opportunity cost"), most airlines have adopted acronyms or other shortcuts to identify the terms. Choose those that are meaningful to you and your staff.

Going back to our simple network, assume the following flight data for a given time period in the SYR–OMA analysis:

1. Local SYR–OMA revenue	$6,000
2. Proration to SYR–OMA of connecting itineraries	$1,500
3. Proration to other segments from connecting itineraries	$4,000

Table 10-3 Three Basic Measures of Profitability

Measurement	*Revenues Included*	*Costs Included*
Variable segment profitability	Local SYR–OMA revenue plus prorate of connecting itineraries using SYR–OMA	Variable flight costs
Variable plus ownership	Same	Variable flight costs plus aircraft ownership costs for the segment
Fully allocated segment profitability	Same	Fully allocated flight costs

Table 10-4 Network Connectivity Components to be Considered

Network Component	*Revenues Included*	*Costs Included*
Basic network contribution	Remaining proration of connecting itineraries using SYR–OMA (not used above)	Variable network costs (connecting costs *plus* variable passenger costs on the connecting segment)
Network contribution with opportunity costs	Same	Same as above *plus* the opportunity cost of SYR connecting passengers to the network

4. Variable operating costs $4,500
5. Aircraft ownership costs $2,000
6. Allocated overhead and non–operating costs $1,500
7. Network variable costs $700
8. Network opportunity costs $500

Now let's review the measures of segment profitability along the spectrum of financial hurdles (the equations below use the numerals assigned to the revenues and costs just listed):

Variable Profitability plus Network Contribution $((1+2+3)-(4+7))$ = $6,300

Variable plus Network with Opportunity Cost $((1+2+3)-(4+7+8))$ = $5,800

Variable plus Ownership, Network Contribution $((1+2+3)-(4+5+7))$ = $4,300

Variable, Ownership, Network plus Opportunity $((1+2+3)-(4+5+7+8))$ = $3,800

Fully Allocated plus Network Contribution $((1+2+3)-(4+5+6+7))$ = $2,800

Fully Allocated, Network plus Opportunity $((1+2+3)-(4+5+6+7+8))$ = $2,300

Variable Segment Profitability $((1+2)-4)$ = $3,000

Variable plus Ownership Profitability $((1+2)-(4+5))$ = $1,000

Fully Allocated Segment Profitability $((1+2)-(4+5+6))$ = ($500)

Does SYR–OMA make $6,300, lose $500, or something in between? The answer depends on who is asking the question and for what purpose. The specific profitability figures we have calculated are not relevant for all decisions. In general, the figure that is most relevant is the one that describes the business problem or opportunity faced by the decision maker at the time. Table 10-5 outlines only a few examples.

With these concepts explained and understood, the remaining discussion will focus on particular issues

Table 10-5 Examples of Business Problems Faced by a Decision Maker

Decision Process	*Relevant Profitability Measure*	*Comments*
Short-term scheduling optimization	Variable with network contribution	In the very short term, ownership and overhead costs are fixed. Flight and market level need the network contribution to be useful.
Middle-term scheduling optimization	Variable plus ownership with network contribution	In the middle term, aircraft may be fungible.
Hub profitability for a single month	Variable profitability, no network contribution	In aggregation, adding network contributions would double-count revenues.
Hub profitability for six months	Variable plus ownership, no network contribution	A combination of the middle-term scheduling and single-month hub profits example.
Hub viability	Fully allocated profitability	Over time, every cost is variable.

related to the measurement and use of the profitability measurement terms defined earlier, as well as a few special profitability measurement topics.

Measuring Revenue

Using revenue in airline profitability measurement is not confusing, and relevant data are generally easy to obtain. Even the simplest revenue accounting processes track revenue at the flight coupon level, allowing for an easy collection of local revenue. Multiple-segment O&D itinerary revenue should also be available from revenue accounting's coupon lift. (This is the physical record that the accounting department uses to "recognize" revenue. Some airlines physically match every single coupon to the passenger's reservation record; others merely use a sample.) Therefore the problem becomes the appropriate proration methodology to use to break down the O&D fare into its segment components.

Depending on the lift process used, the revenue data that are obtainable may not be valid for specific flight-level analysis. A sampling process that is valid only at the market level would not provide accurate flight-level revenue data. As with revenue management, the robustness of the revenue accounting process and the availability of revenue data through this department is a strong determinant of the accuracy and usefulness of the segment profitability exercise.

The proration method chosen must be meaningful in terms of the decisions that are to be made from the resulting profitability data. For example, consider an airline that is principally a hub-and-spoke carrier that wants to use profitability information to determine which routes to include or exclude in the network. In this case, the proration method chosen is less relevant since the meaningful profitability measures will include the network contribution. For airlines requiring or desiring an approach based more on the particular segment, good proration data are needed to capture a meaningful picture of segment-based revenues.

The most common proration methodologies for this purpose are mileage-based and full-fare ratios. Either of these methods is mathematically easy to calculate, and often provide similar results, since full fares typically are better correlated to mileage than are discount fares. Your airline's revenue accounting department is likely already practicing one of these proration methodologies, since proration is necessary to reconcile joint or interline fares.

As with most reporting, airline segment profitability data must be accurately interpreted so that they can be an effective tool for decision making. The revenue proration methodology adopted will affect the ratio of local to network connecting revenue, and this could lead to alternative conclusions on the "stand-

alone" or "network contributor" features of any individual route. Many senior airline managers are uncomfortable when significant capital is tied to routes that are locally unprofitable but that are strong contributors to the network. Short-haul segments with high connecting percentages tend to look this way, since the cost of short-haul service is rarely covered by the proration ascribed to them from longer-haul itineraries.

It may seem that the choice of reporting net versus gross revenues could significantly affect the profitability measurement process. But this is not true, since in all cases the net outcome is identical. The issue of gross versus net revenues should be considered in light of other issues, since it will affect the benchmarking and comparative competitive position of both revenues and costs.

Finally, the question of revenue recapture must be addressed. The profitability analysis is attempting to ask which revenues are at risk if the flight/segment/hub is changed. How much of the specific revenue is truly at risk, given the probability that the airline can recapture this revenue on other flights in the network? For profitability reporting purposes, it is recommended that one ignore revenue recapture since it would tend to overcomplicate the process and raise more questions than answers in the results. This fine level of analytical detail is essential in a schedule and market-planning forecasting tool, and it is here that revenue recapture can be more properly modelled and considered. For profitability reporting, assume reported revenues are 100 percent incremental.

Measuring Variable Cost

Variable flight costs include those costs directly related to flying a specific aircraft between two cities. These are broken into five categories: passenger variable costs, fuel cost, crew cost, direct maintenance costs, and servicing costs.

Passenger Variable Costs

Passenger variable costs are costs associated with having a passenger in a seat versus having an empty seat. The largest of these costs is often commission costs (if gross fares are being used for revenues). Commission costs are usually easy to calculate in general, but to be most accurate it is best to match commissions by route where possible. Once again, the accounting department needs to provide this information accurately, although the sales department should be able to help clarify matters.

Catering is generally the next-largest passenger variable cost, and, like commissions, the closer one comes to pinpointing costs based on actual flight invoices, the better. Fortunately, it is often easy to get this information on a flight-specific basis either from the company's catering group or from the third-party provider, if one exists.

Other passenger variable costs include liability insurance and possibly passenger-specific fuel burn and passenger-specific ticketing fees.

In general, passenger variable costs should be limited to those costs that are actually incurred only when a passenger books and shows up for a flight. If other costs are allowed to be considered under this category, the incremental concept so essential for this exercise is lost.

Fuel Cost

Fuel cost is the cost of the fuel burned on the flight. It should be easy to collect this information on a flight-specific basis since it is typical for fuel to be invoiced by flight.

Crew Costs

Crew costs include the cost of the cockpit and cabin crew assigned to a given flight. These costs are a function of the specific crew contract, yet there are some costs associated with crew that should not be considered in this category, such as crew scheduling, which is more of a general administration expense. Crew costs should include all direct pay and benefits for crew assigned to a given flight, as well as any off-duty or "penalty" pay that is directly driven by that flight's schedule. This measurement exercise can be complex and will be impossible to accurately model without the support of the crew scheduling department. It is important to keep in mind that crew costs are included in the variable cost category, meaning that if the flight does not fly, the costs are not incurred.

Often crew costs are not variable within the bidding window. For a general profitability analysis, you will not need to take this into account as long as the cost is accurately posted on its own. This way, very short-term decision makers can easily add the variable crew cost back as a fixed cost when necessary.

Direct Maintenance Costs

Direct maintenance costs include the allocated cost of maintenance checks attributable to specific operational drivers, such as hours flown or number of landings. Maintenance costs are one of the hardest cost

items to accurately measure at a flight-specific level. One method that is recommended is to work with the maintenance team to break down total direct costs into departure-related and flight-hour-related costs. This way, maintenance unit costs can be applied to each flight activity.

Special care must be taken in figuring direct maintenance costs because most maintenance expenditures are driven by direct flight activity but not charged as such (compare this with fuel or catering, for example). With maintenance costs driving as much as 18 percent of the entire cost structure, fudging in this area is not recommended. A dedicated team from finance and maintenance, led by the maintenance controller if such exists, must address this cost item carefully to ensure that the proper incremental standards are being set to the unit costs, that is, that cancellation of the flight directly results in the saving of maintenance as measured.

Servicing Costs

Servicing costs are the variable costs associated with handling a specific flight. This includes ground handling charges, which are often easier to capture at the flight level when the service is outsourced. If the data are not readily available in a flight-specific format, these costs will need to be broken into components that are variable for the specific flight. The finance team and the airport controller should take the lead in this effort. Airport processing costs are similar; they relate to the ticket counter, gate, and security functions.

Servicing costs create one of the more difficult decisions that needs to be made in order to allocate costs to create an accurate measurement of profitability. Typically, these costs would include only the variable manning and processing charges, and the station overhead and facility expenses would be allocated later in the overhead cost section. But what about a station that has only one flight? When evaluating this scenario, the most accurate variable cost would include 100 percent of the station operating costs, including station-related overhead. After all, if the flight were removed, those costs could also disappear. In practice, however, this is difficult to do, since it requires constant and accurate updating of the total flight activity by station and ascribing the station costs correctly among the flights that arrive at that station. Starting out, station operating costs are a detail that can be avoided but ultimately should be addressed so that one can build the most meaningful decision tool.

Landing fees are straightforward, and the necessary data are likely available directly from the airport invoice, as should be en route navigation charges and any flight-specific security costs.

In general, the variable costs will likely constitute over 60 percent of the airline's cost structure. This set of costs is the single most important factor in correctly understanding the cost side of flight-specific profitability. Over most time periods for which profitability data will be used, the variable costs represent the airline's investment in making that flight work, and these costs can in fact be made to grow or disappear as necessary with proper management.

Measuring Aircraft Ownership Costs

Aircraft ownership costs (which should include hull insurance as well) are generally easy to calculate, but for these costs there is the greatest gap between accounting data and profitability data because of three specific issues: financing, fleet allocation, and flight scheduling.

The actual capital cost for an aircraft is a direct function of the financing method used to acquire the asset. Since the ownership cost is a real piece of the airline's structure, it cannot be ignored. The ownership cost for a particular aircraft may or may not be consistent with the market rate for the aircraft. A perfectly pure capital cost accounting method would always charge the airline the higher of the market lease rate for the aircraft or the actual cost paid. This cost method would not unfairly benefit a route using equipment efficiently leased or owned when in fact the best economic decision may be to lease out the aircraft to another operator at a market rate. This helps ensure that the airline is always optimizing its asset base.

Unfortunately, it would be hard to keep the ownership cost accurate, and efforts to do so will likely be frowned upon by senior management when they see the profitability of a route swinging with aircraft capital cycles (even though it should). It is good to remind ourselves why we want the profitability data: to make accurate decisions about the deployment and use of the airline's assets. Looked at this way, every aircraft should meet this criterion, and thus the costs of ownership or leasing should be charged to the routes it flies on an equivalent market lease rate basis or the actual cost basis, whichever is higher.

Fleet allocation is the next confounding variable in the measurement of aircraft capital in an effort to determine segment profitability. Consider an airline with two 150-seat aircraft options: a fully depreciated 727 and a new A320. If both aircraft could fly a given route, does it make sense to charge one flight the

higher lease rate of the A320 just because someone in scheduling or operations assigned the tail that way? It probably does, because if the variable costs are accurately measured, then the operating efficiencies of the newer aircraft should be modelled for that particular flight. Presumably, the capital cost of the newer aircraft was justified partly because it would create ongoing operating cost savings, and therefore the new aircraft was purchased instead of another used 727.

A theoretically interesting procedure would be to blend both the operating costs and the capital costs of like-size equipment and charge this "blended aircraft" cost to all routes using that size equipment. Thus one would be able to accurately reflect the airline's current fleet mix and not subject individual flights or markets to the seemingly random whims of the aircraft schedulers. An example is an airline currently flying both old DC9s and new A319s. For many missions, either one or the other plane could be substituted, but the operating and capital costs are quite different. It doesn't make sense to have these differences bias the economic view of a market that can be flown by either aircraft. Allowing for such complexity is worth the effort only when a significant piece of the fleet is subject to this type of situation; otherwise this method is not recommended for anyone beginning the effort of measuring flight profitability.

The last and most problematic issue in the area of aircraft ownership costs relates to the flight schedule itself. It at first makes sense to develop for each aircraft an hourly lease rate that ties to the methodology we have already outlined, and then to apply this rate to the actual hours flown on an individual flight segment. But this approach is ill advised, especially for hub-and-spoke carriers, because in order to create the most connections at the hub, often an aircraft flying a shorter-haul segment must wait on the ground at the spoke station until the aircraft flying longer segments have reached the appropriate point in their missions. As a result the actual time flown on the shorter-haul segment often does not include the time that the aircraft is out of service because of other missions. Thus the total time of the rotation—both actual hours flown and waiting time—may be the most relevant measure for this purpose. Consider a flight for which the aircraft flies one hour from the hub, waits on the ground for three hours, and then flies one hour to return to the hub. It may be necessary for the aircraft to wait so that each flight can meet appropriate connections, yet the aircraft is unavailable for use for a full five hours. The ownership costs should reflect this usage.

In summary, aircraft capital costs should be set on a market lease rate basis (or higher, if actual costs are such) and ascribed to individual flights based on the total time commitment of that aircraft to service the flight within the schedule. The most sophisticated users may attempt to model blended variable and ownership costs for equipment of similar size with very different economic characteristics, but this approach would not add significantly to the robustness of the process.

Measuring Overhead and Other Non-Operating Costs

Virtually all other costs of the airline can be grouped together into a single category and ascribed back to individual flights according to some generalized drivers, such as ASMs, revenue passenger miles (RPMs), departures, or block hours. The most important considerations here are to break out each cost separately so that each line item can be tracked and to ensure that all flights aggregated sum to the actual non-direct costs.

The costs included in this category include all corporate overhead, marketing efforts (sales force, advertising, promotions, and so forth), and indirect maintenance and facility charges. Typically not included would be interest and taxes, because the flight profitability database should theoretically tie to the operating line versus net profit line of the company's profit and loss statement.

Some companies may choose not to model overhead and non-operating costs and instead produce a system that measures a flight's contribution to corporate overhead. This method is not recommended, but not purely for reasons of profitability measurement. Rather it is not recommended because it would tend to create an unrealistic view of the airline's profitability among the users of these data, and the indirect effect of this unrealistic view cannot be accurately predicted. Being able to look at fully allocated profitability for every flight in the system is a sobering exercise that tends to break down emotional barriers and slay sacred cows that do not have economic justification.

Measuring Real Network Connectivity Costs

Network connectivity costs, sometimes called "beyond" costs or "upline/downline" costs, represent the incremental cost of adding an additional

connecting passenger into the network from a segment or flight being evaluated. Returning to our previous simple example of the SYR–OMA flight, a passenger flying from SYR to OMA and connecting to LAX would be considered "local" to the SYR–OMA flight and "beyond" or "connecting" to the OMA–LAX flight. The presumption here is that during the evaluation of a unique piece of the network (in this case SYR–OMA), the rest of the network is fixed.

Obviously, this analysis is credible only when we are examining a small piece of the network. In this context, though, it gives a meaningful view of the network contribution of individual flights and markets. Taking the assumption of a fixed network, then, it is logical that the only costs incurred by the network for the additional passenger are the passenger variable costs on the connecting segment. The assumption is that the connecting flight would be flown anyway, so the flight's variable, ownership, and overhead costs are not relevant in this decision.

Consider our SYR–LAX passenger, who pays $250 for the O&D. Our proration methodology gives $100 to the SYR–OMA flight and $150 to the OMA–LAX flight. Based on our previous analysis, the $100 revenue would be applied to the SYR–OMA segment, and it, along with all other local and connecting prorations, would be charged the flight's variable, overhead, and fixed costs, depending on the profitability measure used. The $150 portion of the revenue would be considered "connecting" or "beyond" revenue and charged against it would be only the passenger variable costs of the trip.

Obviously, the margin on the connecting revenue is significantly higher than on the local legs, and this is because we are making the "fixed network" assumption. Again, however, when looking at a finite piece of a complete airline network, fixing the rest of the network becomes a very credible way to evaluate the full network impact of adding another connecting passenger. The only problem with this method of analysis comes in the aggregation of flights, whereby the beyond contributions are double counted. This is discussed later in the section "Aggregating Flight-Level Data into Meaningful Entities."

Measuring Network Opportunity Costs

The piece of the profitability measurement that is most theoretical and most difficult to quantify is the network opportunity costs, sometimes called "beyond displacement" costs. The idea is clear and straightforward. In the previous analysis, we charged the connecting revenue with only the passenger variable costs incurred on the connecting segment. To do so we effectively assumed that the seat on that aircraft would have otherwise been empty. With the network opportunity cost analysis we break down that assumption and reflect the fact that there was in fact some probability that the seat used by the connecting passenger would be put to an alternative use, creating an opportunity cost.

There are at least two methodologies to consider in estimating the opportunity cost, both of which should be understood by the airline's yield management department. The first, and easiest, method is to use a generic "spill table" available from Boeing or Airbus in order to estimate the probability of selling an additional seat given the connecting flight's load factor. The spill figure will give you the probability that the seat filled by the connecting passenger would have been sold. We then apply a marginal fare to that seat in order to create the opportunity cost.

A more sophisticated, and complicated, way to estimate this cost is to look at the rate at which low fares are closed to sale, sometimes called bucket-level closing rates from the yield management system. Although this approach is a bit more complicated, it provides a better estimate of the displacement probability, particularly because the data should be unique to the airline, unlike that provided by the generic spill tables. To avoid confusion, one should not lose sight of the ultimate goal: if we did not fill this seat with a connecting passenger, what is the expected value revenue we would have earned for it?

Let's review a simple example, going back to our SYR–LAX passenger. The connecting revenue of $150 has already been reduced by the passenger variable cost of the OMA–LAX flight. Next we estimate the probability of filling that seat with another passenger, which is a function of the total demand for the OMA–LAX flight. If many passengers are demanding this connecting flight, resulting in a high load factor, then the probability is good that the seat could be filled without the SYR–OMA passenger. In this case, the network opportunity cost would be high, and the SYR–OMA flight would be given relatively little value in terms of network contribution since the connecting passengers are only filling seats that would likely be sold anyway. If the OMA–LAX flight has a lower load factor, then the opportunity cost would be lower, and the $150 from our SYR–LAX passenger would have only a small cost charged against it.

Mathematically it is possible that the connecting contribution could be negative. This would happen if

the opportunity cost of the connecting seat was higher than the prorated connecting revenue. If the yield management and scheduling groups are doing their jobs well, this should not happen, since the O&D would not be sold, or additional capacity would be placed on the clogged connecting leg. It is important to understand that the incremental value of connecting passengers into the network is a direct function of the network's ability to accommodate those passengers economically.

As mentioned earlier, it is always desirable to consider network connecting revenues net of both passenger variable costs and beyond displacement costs. Since the former are easy to calculate and constitute part of the fundamental profitability measurement process, these costs should be calculated right away. The calculation of opportunity costs can be added later, but until such time users of the data should recognize that the connecting benefits are something less than those stated.

Aggregating Flight-Level Data into Meaningful Entities

So far, we have discussed only the case of a single flight and the evaluation of its profitability. For the analysis to be a meaningful decision-making tool, it is necessary to aggregate this information at the market, city, hub, and system levels. Doing so is easy, but some cautions must be taken.

Using the segment-only profitability measures (not including connecting revenues or expenses), you can create the summed profitability for a set of flights by simple addition. The addition of the network contribution makes sense in the flight-to-market consolidation, but not beyond this level, because for the entire network every flight's connecting contribution is counted as another flight's local or segment value. The double counting becomes evident when we try to sum all flights into a hub, for example. While the double-counting effect makes it obvious that this consolidation is not feasible, there is also a methodological reason not to do it: the assumption of a "fixed network" that allows us to charge only passenger variable costs against the revenue breaks down when we need to look at a large piece of the network.

It should be obvious by now that it is possible for the system to lose money (the sum of the P&L for the entire network is negative) but for each individual segment to justify itself based on its contribution with network connectivity. This is a situation that must be looked for carefully, because it is too easy to justify keeping each piece of a losing network since this situation would suggest that the airline couldn't shrink to profitability. The flight profitability information will accurately identify such a situation and aid greatly in evaluating network restructuring options to make the entity profitable.

Proper reporting of the flight profitability data should not even allow sums of figures that include network contributions, since this information could easily fool an unsophisticated user.

Flight Profitability versus Accounting Statements

Earlier we addressed the distinction between accounting statements and flight profitability analysis. It should be clear by now why the flight profitability data will not tie directly to the accounting statements, since the profitability data use costs that are not ledger expenses (such as normalized maintenance, market ownership rates, and network opportunity costs).

Understanding these differences and reconciling the monthly profitability data with the monthly closing of the accounts is a useful exercise. It will help to ensure that the broad areas of overlap are in line and that the assumptions made in the profitability database are clear and understood.

Suggested Reporting Formats

There are many ways to report flight profitability data. Monthly reporting is a logical extension of the accounting process, and the data should be available at the individual flight level if the system has been built according to the process described herein.

The first report to consider lists each flight in the system, subtotalled at the market level (all flights between two cities) by nonstop flight segment. For every flight, the major profitability figures, with and without connecting contribution, should be laid out in order from most variable to a full allocation of the airline's fixed costs. In addition, meaningful flight statistics, such as revenue per available seat mile (RASM), cost per available seat mile (CASM), load factor (LF), stage length, passengers, departures, yield, and block hours, should be appended to each

flight segment since it is logical to see the profitability information in the context of these data.

Next consider a route-by-route P&L, which lays out the revenue and cost items completely and allows for a detailed review of the route's performance for the given month. This report makes the most sense when aggregated at the market segment level as opposed to the level of the individual flight.

The third and most interesting format to suggest is a segment-based trend report, which shows the data outlined at the segment level, subtotalled at the market level by nonstop flight segment, and trended over a twelve-month period. Since many airline markets are seasonal, it is important not to rush to judgment about a particular flight or segment based on a single month of data. Thus the trend report should include a full year's data.

There are undoubtedly many other formats that tie the flight profitability into hub groupings, city groupings, aircraft tail routings, and so forth. Users should get comfortable with the basic data presented as described, and then create alternative formats for specific use based on the airline's individual needs.

Unique Flight Issues and Considerations Concerning Their Appropriate Measurement

The outline described herein represents a description of a useful decision-supporting tool for evaluating profitability at a less-than-complete network level. There are three special cases to discuss wherein the basic profitability analysis may be misleading. These are (1) flights with high frequent flyer redemption rates, (2) flights created with "free" or otherwise unused aircraft time, and (3) flights with multiple stops.

Flights with High Frequent Flyer Redemption Rates

Many airlines have one or two flights in the system that absorb a disproportionate number of frequent flyer reward trips. These flights are likely to show up as highly negative in terms of the segment P&L, since the full costs of each flight will be ascribed against a low passenger revenue base. The true value of these flights will be understated in this case, since the prospect of garnering such flights through the frequent flyer program is presumably at least part of the reason that passengers chose to fly the other segments on the system at the levels they did. A good example is flights from the U.S. mainland to Hawaii, which are typically frequent flyer "loss leaders" for U.S. carriers.

There are at least two ways to approach the problem of high frequent flyer redemption rates for a given flight. The first is to artificially ascribe a revenue figure to each frequent flyer passenger as a function of his or her total value in the entire airline system (other trips, higher average yield, and so forth). With this information, the special flight can be evaluated as usual, since the resulting revenue stream created should at this point cover the flight costs if in fact the flight is profitable to the network. Creating this proxy revenue is a tricky process, but it should be workable with a team including the finance department, revenue accounting, and the frequent flyer group.

A preferred but more detailed option is to create a "mileage earned" and "mileage used" account for each flight in the system and to charge flights that generate mileage liability while crediting flights that relieve this liability. The advantage of this system is that the effect of frequent flyer usage is seen on every flight segment, and since it's a zero-sum game, the total segment consolidation will still tie to a reconciled accounting statement. This process is detailed, but it is useful to the airline with high frequent flyer usage.

Flights That Are Created with "Free" or Unused Aircraft Time

Flights termed "free flying" by the schedulers are flights flown with otherwise unused aircraft time. They are often called "free" because they don't require incremental aircraft capital. Overnight flights or tag-end flights are examples. A fair way to view these flights is by using a profitability measure that excludes aircraft ownership and fixed overhead costs.

On their own, these calculations are not complete, because if certain flights will never be looked at in terms of aircraft ownership or fixed costs, it must be made certain that the total aircraft ownership and fixed costs are ascribed to the rest of the flights in the system. Consider a case involving only two flights: a main flight (for which the aircraft was purchased) and a tag-end flight late in the evening, using the same aircraft. A typical allocation methodology would charge each flight based on its hourly or ASM usage for both aircraft ownership and overhead costs. If the tag-end flight were considered "free" and therefore never evaluated on a fully allocated cost basis, then the result would be that the main segment is

being undercharged in both ownership and overhead categories.

This point is fairly subtle and subject to interpretation, because it is not recommended that an airline maintain a file of "free flights" that would not have charges allocated to them. Care must be taken on a continual basis to justify such routes, and to study the effect on the rest of the system when certain costs are removed and unaccounted for elsewhere.

Flights Involving Multiple Stops

The last special situation is for flights involving multiple stops, or "milk run" flights. For these, the segment-based analysis approach is still valid, but the network contribution component will be confusing and somewhat misleading. In a way, this situation is analogous to the hub or network problem discussed earlier: the entire routing must make sense, even if each hop on the flight is justified because of its "network" connectivity contribution (or, in this case, its next segment) contribution.

In this case it is recommended that one base the profitability analysis principally on the individual flight segment, understanding the portion of revenues that are derived from local O&Ds versus those passengers who traverse multiple stops on the itinerary. The network contribution figures will be helpful in this exercise, but they cannot be accepted without first getting a clear understanding from the revenue accounting department as to how the multiple-stop O&D itinerary revenues have been prorated across the various segments.

Summary

Understanding the profitability of multiple product lines is something every business must do in order to maximize the value of its assets and deploy those assets in the most efficient manner. The airline business benefits from the fact that its principal capital asset, the aircraft, is movable and can be deployed in multiple alternative uses throughout its life. Having meaningful segment P&L information from a well-designed, well-audited profitability reporting process will be invaluable to an airline's decision-making process in the areas of route planning, scheduling, pricing, and fleet planning.

The process outlined in this chapter is a reasonable approach to breaking down profitability analysis of a network into component parts in a way geared toward better decision making. No report or process replaces sound human judgment and interpretation, so it is highly recommended that the airline also invest in talented individuals to ensure that the tools will be used and applied correctly.