Using Risk Assessment for Aviation Demand and Economic Impact Forecasting in the Minneapolis-St. Paul Region

MATTHEW F. HARDISON, RICHARD R. MUDGE, AND DAVID LEWIS

The process of risk assessment was applied to airport strategic planning for analysis of the adequacy of the Minneapolis-St. Paul International Airport. Three steps were used to forecast demand: (a) development of a structure and logic model; (b) development of initial input assumptions; and (c) forecasting, risk analysis, and public exposure. The findings are presented in terms of operational results that define the probability of meeting unconstrained demand under each of three proposed development scenarios. The operational data are translated into implied economic benefits to the region.

Apogee Research and its partner, James F. Hickling Management Consultants, were asked to develop a proposed approach to evaluating the adequacy of the Minneapolis-St. Paul (MSP) International Airport. The existing forecasts, developed for the airport master plan, were correct enough, but the implications were unclear. The planners for the Metropolitan Council of the Twin Cities were faced with the real risk that the forecasts might be wrong. Consequently, the planners needed a process that, in addition to projecting growth in demand, could interpret the forecasts in light of the economic consequences of alternative airport development scenarios.

A request came to Apogee Research, Inc., in the form of a series of questions, the most fundamental being, Are the forecast results correct? Related questions included (a) What are the sensitivities of the forecast to changes in the underlying assumptions? and (b) What do each of the development alternatives imply for long-term development in the region? As a practical matter, any technical evaluation also had to bring together the diverse and divisive groups involved in airport planning if the process were to prove successful.

Simply reviewing the existing forecast would add no new information and was unlikely to create the consensus that would be necessary for long-term investment decision making. Thus, Apogee began by asking a new question: What is the risk that the forecast will be wrong?

If the forecast were framed in terms of the probabilities of meeting demand under selected airport development (i.e., capacity) scenarios, the impact (economic or financial) of different decisions on the region could be evaluated. However, to do so would imply the simulation of a wide variety of inputs for the entire planning horizon—in this case through the year 2018.

The process selected to carry out the simulation was risk assessment. The risk assessment framework, customized for use in aviation strategic planning, would simultaneously evaluate the potential variability in the forecast inputs and therefore the potential variability in the outputs.

RISK ASSESSMENT

Forecasts are often used to make major long-term investment decisions, as they should be. However, forecasts are generally wrong. Although a forecast may, for example, correctly pick the direction of change (i.e., growth or decline), the magnitude of the actual change is often far different from that forecast.

For example, two key parts of the technical portion of an airport analysis—the expected amount of air traffic and the future capacity of the airport—require long-term forecasts of economic, social, technical, and political factors. In most long-term planning efforts, these forces are rolled into a single discrete forecast or a set of discrete forecasts. Each of these forecasts results from a series of explicit and implicit decisions about the many variables that influence the forecast. Some may be highly unlikely, such as explosive traffic growth with no capacity improvements. Although this traditional methodology of probable or expected outcome helps focus the decision-making process, it provides no guidance regarding the likelihood of a given outcome, thereby leaving the community and its elected decision makers with an incomplete view of the future.

Hence, forecasts themselves often become a major focus of local debate. Those in favor of a given solution will, of necessity, rely on projections in justifying their proposed approach. However, those in disagreement with the forecasts will offer resistance, leading to a protracted debate among experts. Moreover, those opposed to the approach for other reasons (such as its implications for the pattern of economic development) will also focus on the projections, knowing full well that virtually every important assumption underlying a projection will, to some extent, be wrong. In this way, opponents of a given investment can pose a serious and effective threat to the planning process by creating plausible scenarios that differ from those underlying official projections.

Should decision makers and planners ignore forecasts? Of course not. Public participation and debate over forecasts is

M. F. Hardison and R. R. Mudge, Apogee Research, Inc., 4350 East West Highway, Suite 600, Bethesda, Md. 20814. D. Lewis, James F. Hickling Management Consultants, Ltd., 350 Sparks St., Ottawa, Ontario, Canada K1R758.

natural and productive, and should be encouraged. But how can forecasts be developed and presented in a way that defuses the unproductive acrimony and manipulation that so often plague forecast-related planning efforts? Reliance on high and low cases has proven to be of little value in this regard because they, like the point-estimate forecasts, indicate nothing about the relative likelihood of any given outcome. Perhaps even worse, they are usually developed by assuming all variables change in the same direction—an outcome that is just as unlikely as all assumptions being accurate.

Probability provides a way around the limitations of the discrete point-estimate forecasts by describing the confidence, or odds, that an expected outcome will actually materialize. To understand how probability aids decision making, consider a simple example. Before the advent of powerful computers, weather forecasters would simply assert their mean expectations: "We do not expect rain today." The decision on whether or not to hold a picnic would be easy. Now, the same forecast incorporates the probability for each causal factor in the determination of rain, and the forecaster announces, "There is a 25 percent chance of rain by midafternoon." A more reasoned decision regarding the picnic is now possible. If the event involves costly logistics for hundreds of people, a rain date might well be announced. In the past, provision for risk was not possible, and many dollars—not to mention goodwill and tempers—were lost.

A similar process—one that would integrate probability with the existing forecasting methodology—would clearly assist in resolving many of the key questions posed by Minneapolis in particular, and those raised during airport strategic planning in general. The approach adopted for the Minneapolis-St. Paul analysis, termed "risk assessment," is based on risk analysis techniques strongly grounded in statistical theory. By quantifying the risks of each of the key inputs to a forecast, risk assessment allows explicit recognition of those factors that are only implied in traditional estimates. Instead of the point-estimate results generated by most air traffic forecasts, for example, the process yields a probability distribution around each key output that more accurately portrays potential variability over time. As such, this tool allows flexibility in policy development by documenting the trade-offs of different levels of service and the ability to plan for a full range of outcomes.

DEVELOPING THE FORECAST

Forecasting demand is the critical first step in the strategic planning process because forecasts serve as the basis for all strategic planning decisions: determining the expected adequacy and longevity of the current facilities, the cost of development alternatives, the implicit quality of service the region provides, and the potential economic benefits of the alternatives. MSP International Airport already had point-estimate demand forecasts in place. These forecasts were used as the basis for modeling, both to provide consistency with existing results (the expected values of the forecasts would approximate the point-estimate forecasts of the existing forecasts) and to ensure that the underlying demand model did not itself become the point of debate. This was accomplished in three steps:

- Development of a structure and logic model;
- Development of initial input assumptions; and
- Forecasting, risk analysis, and public exposure.

Development of Structure and Logic Model

The first step was to become thoroughly familiar with the methodology of the existing analyses, including forecast demand and capacity estimates. This step was not merely a review, but involved the development of detailed structure and logic diagrams for the entire forecasting process and for each category of traffic (such as air carrier, air freight, and general aviation). In this effort, maximum use was made of the forecast framework developed for the region's master plan to ensure direct comparability of results.

Figure 1 shows a structure and logic diagram of the forecasting process for air carrier traffic that was developed for MSP International Airport. Similar diagrams were developed for the remaining categories of traffic. The purpose of these diagrams was threefold:

- To document precisely how the different assumptions of the existing forecast are combined to produce forecasts of each of the key input variables (annual and for various times of day and year) for each category of traffic;
- To obtain agreement from the client and interest groups that the process was properly understood and specified (the structure and logic diagrams then became the basis for algorithms programmed into the risk analysis software); and
- To present the forecasting procedure in a way that the public and interested parties could understand.

The process shown in Figure 1 identifies both the econometric and the accounting (or nonstochastic) and judgmental aspects of the forecasting process. The econometric relationships were fully exposed, along with ancillary assumptions, in the subsequent step.

Development of Initial Input Assumptions

Once the essential data requirements were in place, all input assumptions to the risk assessment process had to be developed. In the case of forecasts of demand and financial feasibility, these assumptions included specific demographic and economic variables for the region that were in turn used to develop the activity forecasts.

Baseline input assumptions (e.g., population, employment, fare elasticity, and various aircraft variables) were established by drawing from the airport master plan (for consistency). The risk analysis approach requires that a probability distribution be attached to each input assumption. Statistical analysis and judgmental factors were used to assign the initial probability distributions. However, to facilitate community and outside expert involvement, expert panels were assembled that were responsible—after a briefing on the technical approach and participant responsibilities—for confirming or adjusting the initial distributions, as necessary. The process had two key benefits:

- Those groups or individuals with a special interest (such as airline, business, or community leaders) were part of the

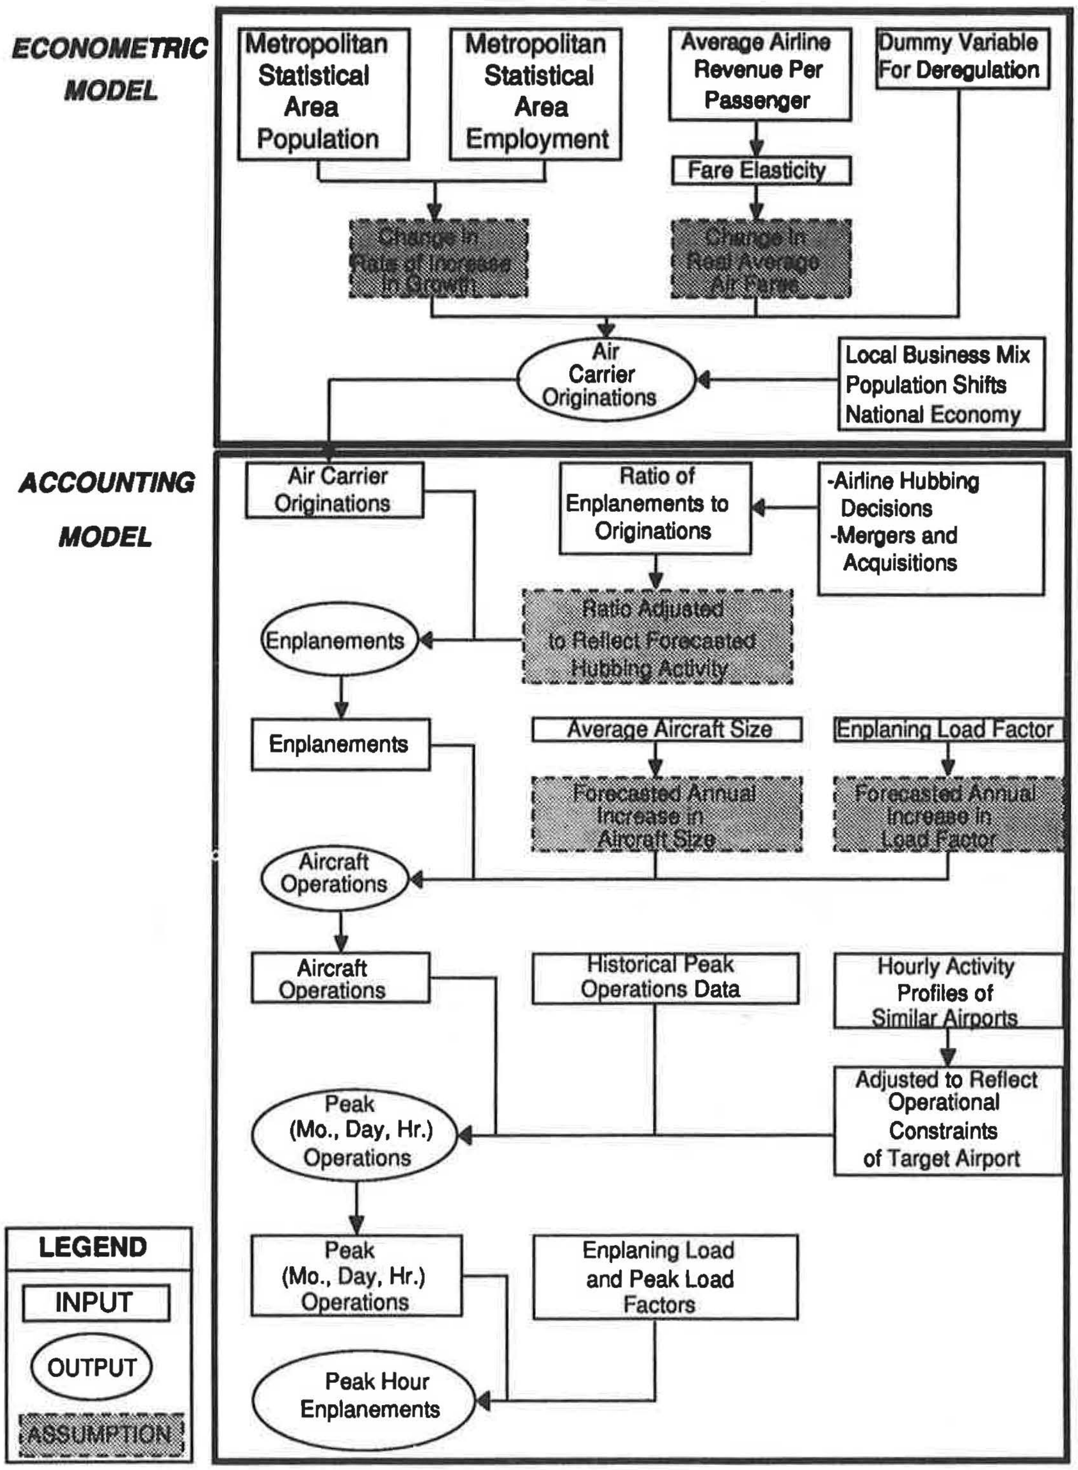

FIGURE 1 Logic diagram for forecasting model (domestic air carrier activity).

process of developing assumptions and thus were involved before actual traffic forecasts were developed (thereby eliminating a great deal of unrest over the results later on); and

- The process facilitated the application of full-scale risk analysis by identifying the estimated probabilities associated with all inputs.

Forecasting, Risk Analysis, and Public Exposure

Once the input assumptions and ranges were developed, computer software translated the ranges into formal probability distributions (called probability density functions). With these in place, the computer software used the forecasting methodology (based on the structure and logic flows programmed into the software) to generate traffic forecasts. The forecasts were developed using Monte Carlo simulation, in which the computer calculates each forecast an unlimited number of times (generally 1,000) by sampling randomly from the various probability distributions. Thus, instead of a point estimate of traffic for each forecast year, this process generated a probability distribution (see Figure 2).

The mean of the distribution often corresponds closely to the point estimate that the traditional process would yield. This relationship is important to understanding the interaction between the traditional forecasting technique and risk analysis.

Once this point in the analysis was reached, changes in the underlying assumptions had little effect on fundamental results.

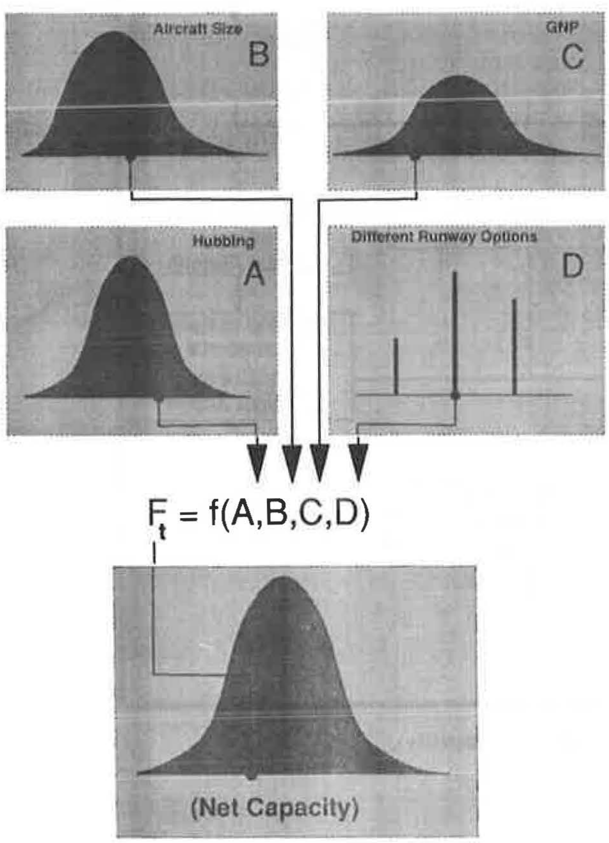

FIGURE 2 Monte Carlo simulation: a way to combine probabilities.

For example, an 85 percent probability of exceeding a given number of operations might change to an 80 percent probability with a fairly large change in key assumptions. This approach allows those parties with a special interest in the results (such as a community or regional planning agency) to quickly put into context the significance of debate over the effect of changes in the model inputs. In this way, special panels or workshops held to discuss the findings defused potentially unproductive debate and gave the forecasts the credibility needed to support the subsequent planning process.

FINDINGS

Forecast results are presented in two ways. First, operational results are provided that define the probability of meeting unconstrained demand under each of three proposed development scenarios. Second, the operational data are translated into implied economic benefits to the region.

Operational Analysis

Although there are different ways to measure the level of service, the approach used in this study was based on the probability of being able to meet expected future demand. A low probability will result in a low quality of service. A high probability, however, runs the risk of overinvesting or building too soon if demand does not materialize.

Traditional planning efforts implicitly assume only a 50 percent chance of meeting expected demand (or, conversely, a 50 percent chance of having adequate capacity). Providing a higher level of assurance would imply a higher quality service to air travelers (fewer delays and shorter delays) but would also require higher costs (financial as well as political) and be likely to increase environmental and other negative effects on the surrounding community.

Figure 3 combines the baseline forecasts with three capacity options contemplated at MSP International Airport: Strategy A, Strategy B, and Strategy C. The base case included only those capacity options already programmed. Strategy A included base case improvements plus a new north-south runway. Strategy C comprised Strategy A and a third parallel.

The results reflect not only each of the alternative runway layouts, but also the likelihood of future air traffic control improvements and the full range of variables that affect future demand. Figure 3 uses expected instrument flight rules (IFR) capacity estimates for each option because the focus is on the year 2008; however, the difference between these estimates and the visual flight rules (VFR) was small.

Net capacity is shown along the bottom of the figure. Negative numbers indicate a shortfall in capacity, whereas positive numbers show capacity in excess of the expected demand. The vertical scale indicates the probability that expected demand will be less than the capacity shown, that is, that the airport will be large enough. For example, there is only about a 5 percent chance that the base case will provide adequate capacity in 2008. However, there is a 23 percent chance that Strategy A will be adequate and a more than 85 percent chance

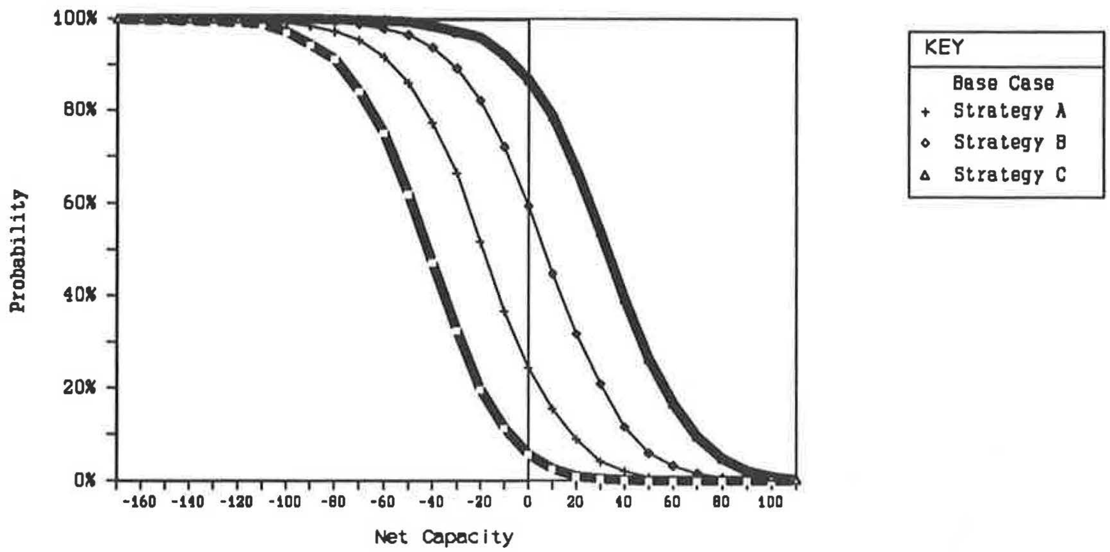

FIGURE 3 Net capacity probabilities for MSP International Airport in year 2008.

for Strategy B. By 2018, as presented in Table 1, the choices are more limited. In that year, Strategy A offers little chance of being able to meet expected demand, although Strategy B shows an almost 45 percent probability of meeting demand. The analysis further indicates that Strategy B represents close to the maximum operating capacity for the current site.

These findings make clear the operational implications of each of three airport development strategies, and, as was suggested earlier, the question of the accuracy of the forecasts now needs no answer. However, from the planner's perspective, a new question is raised: What are the costs and benefits of each of the alternatives?

The technical analysis evaluated both engineering costs and direct, quantifiable benefits (such as fuel savings to the airlines as a result of improved capacity). Financially, this analysis demonstrated that all capacity improvements would realize a high rate of return.

Economic Benefits

Airports often play a key role in attracting new business and in encouraging existing businesses to expand. Economic development was a motivating factor in the decision to build a new Dallas Airport in the early 1970s and in the current plans to build a new airport in Denver. However, strategic planners often make key investment decisions in the absence of any understanding of the potential economic benefits (or costs) of each alternative. Consequently, an analysis of the range of potential economic benefits associated with the development scenarios was prepared on the basis of the risk assessment results. The findings of the MSP International Airport analysis provided information on the range of benefits attributable to aviation under each growth or development scenario and made clear the value of committing to continued growth.

TABLE 1 PROBABILITY OF ADEQUATE CAPACITY IN MEETING PEAK-LOAD DEMAND

	Year		
	1998	2008	2018
Option	Percent		
Base case	10	5	1
Strategy A	67	18	9
Strategy B	92	87	40

Airport economic impacts are of two types: (a) direct impacts related to handling and servicing aircraft, passengers, and cargo, and (b) indirect impacts as these streams of activity move through the economy. The summary numbers presented in this section combine both types of impacts to emphasize the general findings—the range of potential impacts associated with airport development scenarios.

The potential regional economic gains (direct and indirect) that could be achieved from increasing capacity to meet expected demand are quite large (see Figure 4). It is estimated that MSP International Airport currently contributes more than $2 billion a year to the region's economy. On the basis of the airport economic impact model developed for the Minneapolis Chamber of Commerce, the unconstrained forecast of enplanements and aircraft operations suggests that this impact could increase to about $3.5 billion over the next two decades.

However, if a significant increase in capacity is not realized, the severe capacity constraints forecast for the base case development scenario imply annual losses of $1 billion or more by 2008—a loss of one-third. This estimate assumes that the overall level of activity through MSP International Airport will drop in line with the expected higher delays. In other words, rather than impose huge delays on air travelers, the airlines will shift activity to other airports by early in the next century. Clearly, even without significant capacity improvements at MSP International Airport, activity would still increase regional income and add new jobs. However, although the regional economy would not shrink if significant capacity additions were not made, future growth would be limited.

In addition to the value of economic activity, an airport also brings new jobs. One that is capacity constrained, on the other hand, will limit growth below its potential. As dem-

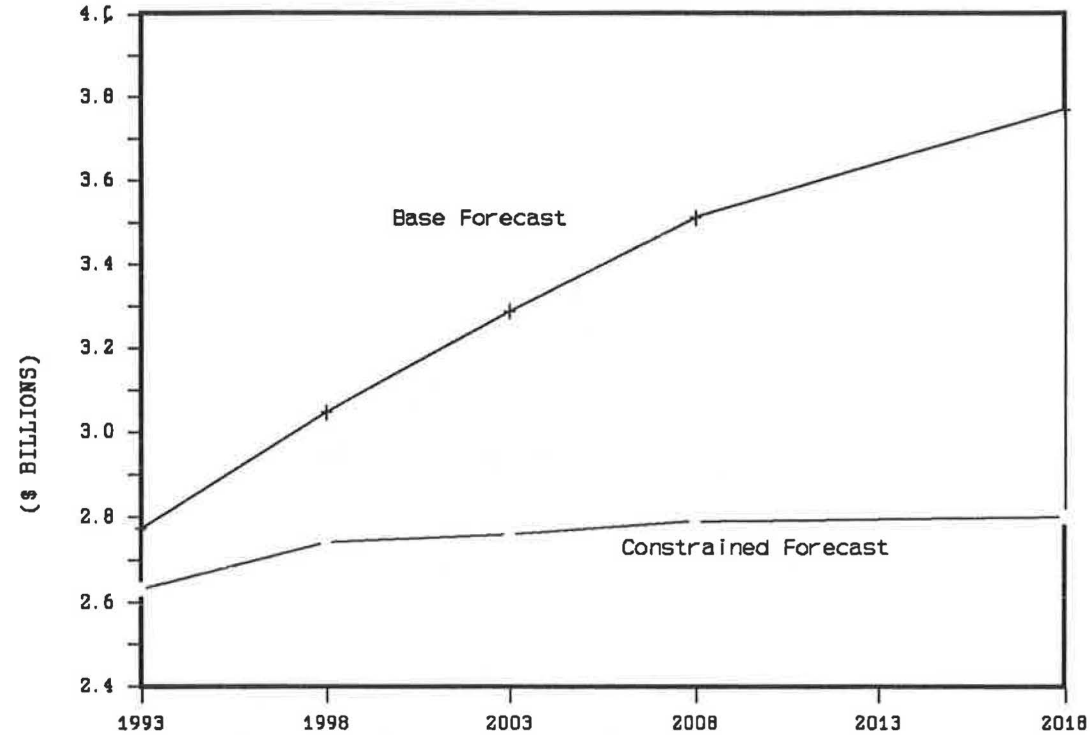

FIGURE 4 Range of potential economic impacts attributable to alternative airport demand scenarios at MSP International Airport.

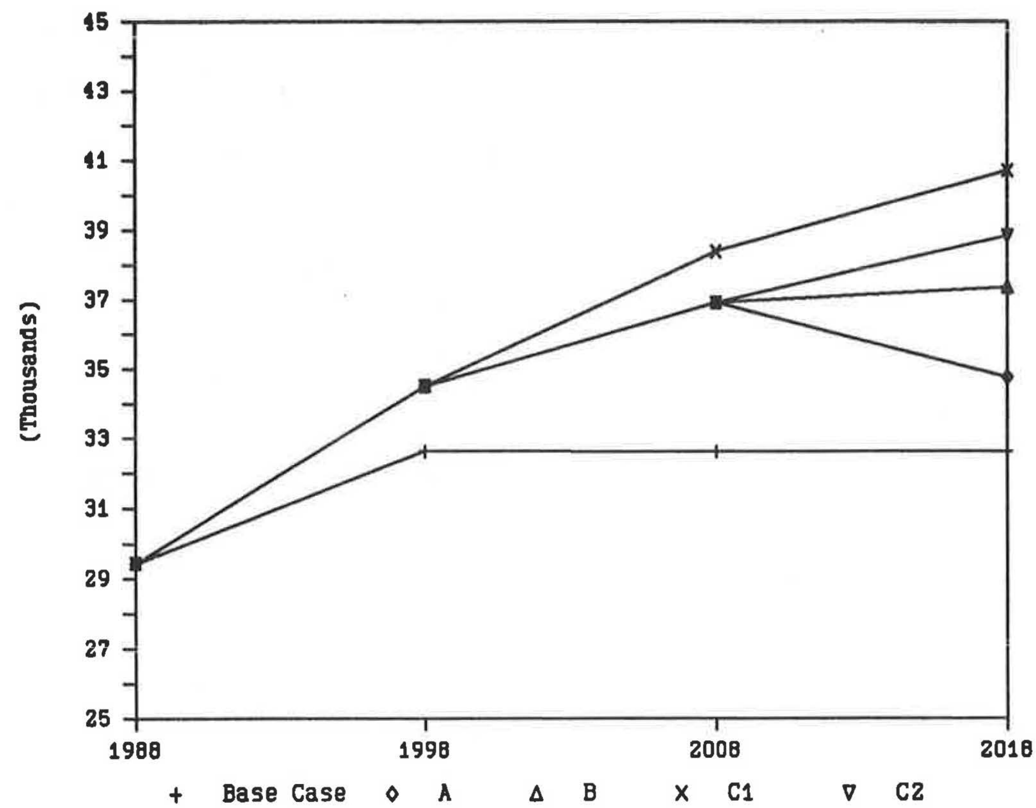

FIGURE 5 Range of potential employment impacts by airport development scenario.

onstrated in Figure 5, for example, the change in service implied by increased delays can only result in lost opportunities for growth. Here, the greater delays and diverted flights under the base case result in a loss of thousands of jobs in the region.

These estimates are based on changes that either restrict capacity to well below current demand or allow capacity to grow roughly in line with expected demand. Yet even the high end of the range of potential impacts may understate the potential growth; by offering a higher level of service, an airport with excess capacity could act as a "growth pole."

CONCLUSION

The risk assessment process has been applied successfully in a variety of analyses. Until the Minneapolis analysis, however, the only related application was an evaluation of the costs and benefits of the nation's air traffic control system. Since its introduction to airport strategic planning, the process has become accepted not simply as a means of identifying the implications of operational, financial, and economic forecasts, but also as a particularly useful tool for redirecting the debate away from the forecasts and toward a decision.

Did the process work in Minneapolis? Before the analysis began, the general decision taken from the airport master plan had been that no decision was necessary—capacity was adequate. And it would have been for the next 5 to 10 years. However, on release of the analysis, the Minnesota state legislature mandated that a large tract of land being considered for sale be held until a decision could be made on how to meet the capacity needs of the region. At the same time, the Metropolitan Council of the Twin Cities began to explore extensive capacity additions at the existing site as well as the possibility of a new airport. The analysis had taken the original forecasts and explicitly identified their strategic, long-term implications. As a result, the regional planning agency and special interest groups together decided to pursue additional economic opportunities rather than face regional opportunity costs. The process had worked.

Beyond these results, the process of preparing the analysis led to the identification of new opportunities for the application of risk assessment. For example, the economic impact of each development scenario depends on the baseline used and assumptions pertaining to the redistribution of aircraft demand at capacity-constrained airports. The analysis presented here focused on comparisons with what would happen under the base case with a predefined set of changes in aircraft mix. However, the base case is itself a moving target, with the practical capacity of the airport continually changing. In addition, faced with severe capacity constraints, the aircraft mix will likely either change as a result of fees (as was proposed at Boston recently), through restrictions, or naturally (as often happens when pilots of general aviation aircraft find capacity-constrained airports less attractive). Recent work with Martin O'Connell Associates, a firm specializing in aviation economics, points to the ability to define direct employment implications even to the level of aircraft mix scenarios. Such results would prove invaluable documentation for an informed and defensible decision on the priority to give different types of aircraft operations.

The process has proven to be a practical, politically sensitive approach to strategic planning and analysis. Its degree of precision can be tuned to meet specific strategic planning needs because the focus is no longer on the forecasts, but instead on the risks.