Aviation System Consulting, LLC
805 Colusa Avenue
Berkeley, CA 94707-1838

PDX Airport Futures Project

Forecast Peer Review

STATUS REPORT

April 15, 2008

Summary

This report presents preliminary comments on the forecast methodology and forecast results prepared to date by the Port of Portland's aviation consultant team for the Portland International Airport (PDX) *Airport Futures* project, based on analysis undertaken as part of the Forecast Peer Review. The objective of the peer review is to provide the City of Portland, the Port of Portland, and the PDX Airport Futures Planning Advisory Group (PAG) and its Forecast Subcommittee with an independent assessment of the forecast methodology, forecast scenarios, and associated assumptions, including the treatment of uncertainty in the forecast process.

At the start of the forecast development process it was recognized that it would be much more productive to have the Peer Review Consultant interact closely with the Port's Consultant (Jacobs Consultancy) during the preparation of the forecasts rather than commenting after the fact. This would allow any concerns to be raised at an early stage in the process while there is time to address them and provide an opportunity for the Peer Review Consultant to make suggestions as the work is proceeding. As a result, the Peer Review Consultant has been closely involved in the forecast development throughout the process, has been included in the distribution of draft products and working materials, including technical memoranda on the model estimation and drafts of presentations to the Forecast Subcommittee and PAG, and provided the opportunity to comment prior to their presentation. Many of the comments provided to Jacobs Consultancy, particularly suggestions on presentations to the Forecast Subcommittee and PAG, have been incorporated into the final presentations. Overall, this has been a very productive relationship.

A central feature of the process adopted for the current forecast is the decision to develop a probabilistic approach to generating the forecasts, also termed a risk analysis. This represents a significant technical improvement over the typical approach adopted in most airport master plan forecasts, including the previous master plan forecast for PDX. The principal advantages of this approach is that it provides a much better representation of the uncertainties inherent in forecasting future traffic at an airport and provides a basis for assessing the robustness of alternative facility plans considered in the master plan to different future traffic levels. In particular, it allows a more quantified approach to understanding the risks involved if traffic grows faster or slower than expected.

However, an important consequence of the decision to adopt this approach is that the analysis is significantly more complex than is usual for an airport master planning study and

involves many difficult technical issues. This was further complicated by the recognition of the need to give detailed attention to assumptions regarding the future price of oil and the impact of policies to address climate change, aspects that have generally been ignored in past forecasts. As in other areas of transportation planning, the Portland region is well ahead of the state of the art elsewhere and is setting new standards of practice for aviation forecasting. This is something that the PDX Aviation Futures project and the Port's aviation consultant team can be justifiably proud of. However, the price to be paid for this capability is that the process is technically challenging, complex, time-consuming and costly. Many of the analytical steps have had to be developed specifically for this study and assumptions made on the basis of limited information.

This report addresses three broad aspects of the forecasting process and is intended to help the PAG understand some of the implications of the technical approach adopted for the forecast, the assumptions used in the analysis, and the range of forecast values generated by the probabilistic modeling. This summary provides an overview of the principal issues identified in the review. More detail is provided in the following sections.

Forecast Models and Probabilistic Forecasting Approach

The forecast models for both enplaned passengers and total air cargo are aggregate models that explain the total traffic at the airport in terms of regional values of the causal variables. This is a fairly standard approach in airport master planning forecasts, but as with any modeling process, the ability of such models to predict future traffic is critically dependent on the choice of variables and structure of the model. A model can fit the historical data quite well and yet produce very inaccurate forecasts, as illustrated by the models developed for the last PDX master plan update. In deciding whether a particular model is likely to produce reasonable forecasts, it is important to consider the variables included in the model, the structure of the model, and estimated values of the model coefficients. Each of these aspects has implications for the ability of the model to generate reasonable forecasts.

In considering whether the current models used in the forecast process are likely to produce reasonable forecasts, the peer review has identified the following issues:

- The use of an aggregate model for forecasting implies that past trends in the composition of the traffic across different markets will continue in the future.
- The use of a constant dummy variable in the enplaned passenger model to account for the drop in passenger traffic since 2001 implies that the proportional reduction in traffic since 2001 will continue in the future. If this effect does not continue, or significantly reduces, the model would under-predict the traffic by as much as 30%.
- The elasticity of demand with respect to both population and income in the air passenger model are lower than would be expected from a general understanding of the determinants of air travel. If the demand elasticity is in fact higher in the future than the model suggests, this would cause the model to tend to under-predict future traffic growth. However, any such effect is likely to be relatively small compared to the overall traffic growth.

The first two issues are intrinsic aspects of the passenger model that the PAG should be aware of in deciding whether to accept the current forecasts. It would of course be possible to generate different forecasts with some assumed reduction in the post-2001 effect. In order to understand the implications of the third issue for the passenger forecasts, it would be desirable to conduct further analysis on the sensitivity of the forecasts to changes in the implied demand elasticity.

The latest air cargo model is based on an underlying logistic growth relationship in the total weight of air cargo per thousand dollars of regional personal income (expressed in constant dollars) that the model development team has termed the cargo intensity. This relationship is adjusted by a term that reflects changes in the price of oil. In addition the model includes a dummy variable that reduces the predicted cargo tonnage by 28% for the years after 1999. The demand elasticity with respect to oil price is about –0.16, which appears reasonable from a consideration of the proportion of integrated air cargo carrier operating costs accounted for by fuel.

The forecasts of future air cargo traffic are significantly affected by the continued increase in the logistic growth relationship, which implies that the air cargo intensity will reach about 37% above 2006 levels by 2035. In deciding whether to accept the current forecasts, the PAG should consider whether this appears reasonable. The PAG should also decide whether the reduction in cargo intensity since 1999 is likely to continue in the future and by the same percentage.

The probabilistic forecasts use a Monte Carlo simulation approach that is highly dependent on the assumed distributions of the explanatory variables in the passenger and cargo models. The effect of these variables has been assumed to be independent, although in reality there is likely to be some correlation. Assuming the variables are independent will tend to reduce the spread of the forecasts. It would also be desirable to perform some additional analysis to quantify the likely impact of this effect.

In order to translate the forecasts of passengers and cargo into forecasts of aircraft operations, a process was developed to allocate the forecast traffic to individual markets (or market segments in the case of air cargo) and then estimate the number of aircraft operations needed to serve each market. This process considered the airlines currently serving each market, the aircraft fleet that they currently operate and any anticipated future changes in their fleet, and trends in average load factors. The process implicitly assumed that the basic structure of the airline industry remains unchanged from 2007 and in fact is independent of changes in the determinants of future demand, particularly the price of oil. While accounting for these effects at the market level may be beyond the resources of the current project, the limitations that this imposes on the resulting forecasts of aircraft operations should be understood.

Forecast Assumptions

Forecast values of future air traffic are only as good as the assumptions on which they are based. The review identified a number of key assumptions that should be borne in mind in considering the forecasts and that appear deserving of further analysis to clarify their impact on the forecasts:

- The use of current U.S. Department of Energy (DOE) projections of the future price of oil. In recent years the DOE has significantly under-predicted the increase in the real price of oil.

- The assumed price of carbon offsets for aircraft emissions are based on carbon offset price levels identified in a recent study by the Massachusetts Institute of Technology, but are reduced below those levels by assuming that only a proportion will be passed on the airlines. They also do not appear to include a factor to account for the effect of radiative forcing.
- The assumed distributions used in the Monte Carlo simulation reduce the average per capita income levels below that forecast by Portland Metro and limit the range considered for the price of oil to the low and high scenario of the DOE oil price forecasts.

Use of Probabilistic Forecasts

The value of the probabilistic forecasts, apart from being more transparent about how uncertain the forecasts are, lies in the ability to address questions such as "how soon might we need the capability to handle 30 million annual passengers at PDX?" or "what is the likelihood that we will need to build another terminal concourse by 2015?" One useful way to use the graphs that show the varying percentiles of forecast traffic level for any given year is to take a horizontal line at any given traffic level and express the percentiles in terms of the corresponding future year at which the given traffic level will be reached.

Conclusion

The various issues identified in this review suggest that a number of the steps in the forecast process could benefit from further study and development, and that further examination of some of the assumptions appears warranted. At a minimum, this additional information will increase the confidence in the forecasts, even if it is decided that there is no reason to change the current forecasts. On the other hand it may suggest aspects where it would be appropriate to make some revisions to the forecasts. In either case, the sooner that this is known, the better.

While this additional work need not delay starting the remaining steps in the master planning process on the basis of the current forecasts, it would be in the Port's and City's interests to continue to refine the forecasting process in parallel with this other work, so that there is a better understanding of the sensitivity of the current forecasts to the issues identified in the review as the other steps in the master plan evolve. It would also be desirable to continue to develop the forecasting methodology so that future updates of the forecasts can benefit from better models and analytical tools. This will be particularly helpful when the forecasts are revisited closer to the time when important facility development decisions need to be made.

Particular issues that appear deserving of further work include:

- Improving the passenger and cargo forecast models to better account for recent changes in the airline industry and reduce the importance of dummy variables in predicting recent traffic trends.
- Development of a market-based forecast model so that future changes at a market level can be explicitly forecast rather than assumed.
- Development of a more detailed representation of airline costs.

- Obtaining more detailed regional cargo data.
- Further study of likely impact of climate change policies or public attitudes on the demand for air travel.
- Further study of the most appropriate probability distributions for the Monte Carlo simulation and consideration of correlation between future values of explanatory variable.
- Development of analytical software to simplify the process of performing the forecast analysis with different assumptions.

The remainder of this report provides a more detailed explanation of the different issues discussed above. The next section discusses the forecast models that have been developed for enplaned air passengers and total air cargo. The following two sections review the methodology used to develop the probabilistic forecasts and derive the forecasts of aircraft operations. The subsequent section discusses the input assumptions used to generate the forecasts. The last two sections discuss the use of probabilistic forecasts in airport master planning and suggest areas where further work could improve the usefulness and robustness of the forecasting process.

Forecast Models

The forecast models for both enplaned passengers and total air cargo are aggregate models that explain the total traffic at the airport in terms of regional values of the causal variables. They are also both multiplicative models (linear in logarithms), which has the consequence that the model coefficient for each variable can be interpreted as the direct elasticity of demand with respect to that variable.

Of course, in reality the total traffic is composed of the traffic in many different markets (e.g. different flight destinations) each with different prices and service levels. A given percent change in price is likely to affect traffic in each market differently. Similarly, a given change in the average price across all the markets could be the result of very different changes in each market. An aggregate model thus attempts to measure the combined effect of all these changes in terms of the change in the average values of the causal variables. In the case of variables defined at a regional or national level, such as population or the price of oil, a given change applies to all the markets equally, although that change may affect each market differently. For example, an increase in the price of oil may affect some markets more than others.

It follows from this that the use of an aggregate model for forecasting implies that past trends in the composition of the traffic across different markets will continue in the future. This needs to be borne in mind in interpreting the forecasts produced using such models.

Passenger Model

The passenger model explains total enplaned passengers at PDX in terms of three continuous variables (regional population, regional average per capita personal income, and average airline yield) and three dummy variables. The airline yield variable (defined as the average airfare revenue per passenger-mile) uses the average PDX yield for years from 1990 and the average U.S. domestic yield for years before 1990. One of the dummy variables is used to adjust for the use of a different yield variable prior to 1990. The other two dummy variables attempt to account for the drop in traffic from 2001. One dummy variable is used for 2001 and the second is used for 2002 on.

Figure 1 shows the fit of the model to the historical data. It can be seen that the model fits the actual data fairly well until 2000, although it under-predicts the traffic in 1996 and then predicts a higher growth rate from 1996 to 2000 than actually occurred, resulting in an over-prediction in 2000. Without the effect of the two dummy variables for the years from 2001, the model would have significantly over-predicted the traffic. The dummy variable for 2001 forces the model to predict the actual traffic in that year. The model then slightly over-predicts the traffic in the next three years, but under-predicts the recovery since 2004. This is to be expected since the effect of the dummy variable from 2002 provides the same percentage reduction for each year. To the extent that the drop in traffic following 9/11 and the SARS scare were most pronounced in 2002 and 2003, a constant adjustment would tend to under-predict the reduction in those years and over-predict the reduction as the recovery proceeds. This has two important implications for the use of the model for forecasting future traffic. The first is that applying a constant proportional adjustment for the years after 2002 is likely to under-predict any continuing recovery of traffic levels from the reductions in 2001 and 2002. The second is that without the adjustments from the dummy variable for years after 2001, the model would greatly over-predict the traffic levels. It is hard to believe that five years after 9/11 the lingering effects have reduced the traffic by almost a third. What is much more likely is that a model based on the relationships between traffic and the causal variables during the era of rapid traffic growth from 1990 to 2000 is not providing a good representation of these relationships in the current environment. The implications of this are discussed further below.

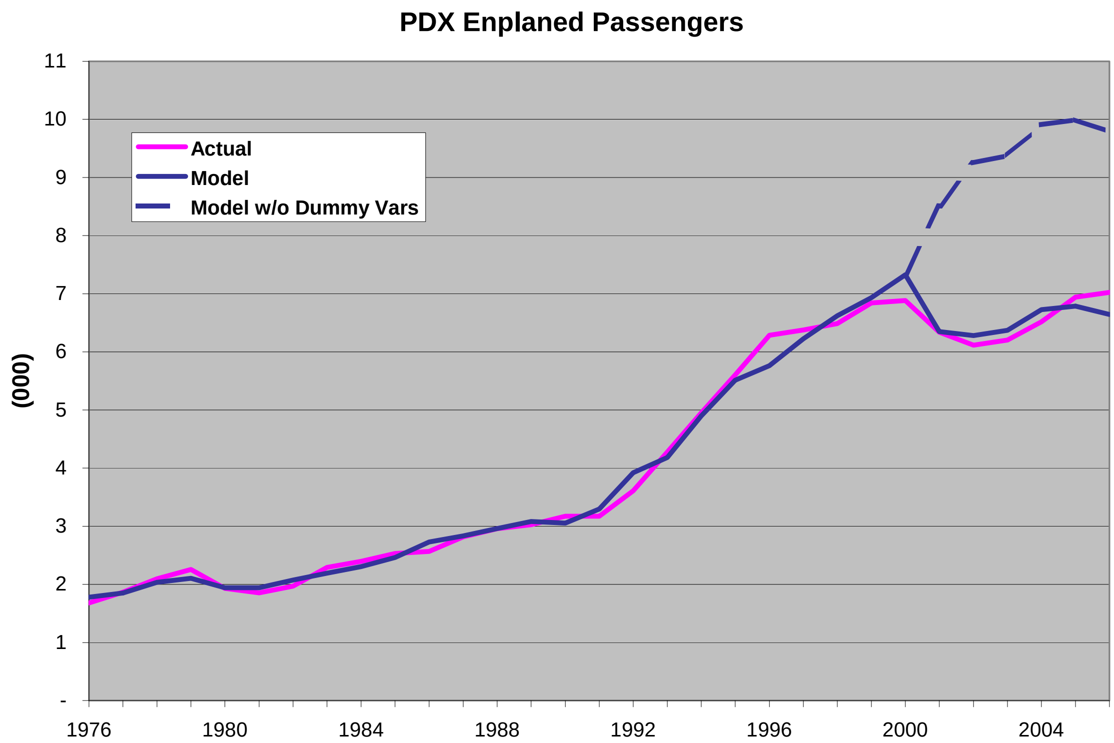

Figure 1 Enplaned Passenger Model

In considering the fit of the passenger model to the historical data, it should be borne in mind that the total enplaned passengers at PDX consist of four different market segments:

1. Originating passengers from the Portland metropolitan area
2. Originating passengers from outside the Portland metropolitan area who travel to PDX by ground transportation
3. Visitors to the Portland region (using ground transportation to access PDX from within and outside the metropolitan area)
4. Passengers connecting between flights at PDX.

Use of Portland metro area population and income and PDX average yield in the model implicitly assumes that the traffic in market segments (2) through (4) maintains a consistent relationship to the traffic in segment (1). Market segment (2) is likely to be fairly small but the proportion of connecting passengers at PDX was around 15 percent in 2005. Both segments can be expected to decline as a percent of PDX traffic as air service improves at other airports in Oregon. Changes in these shares over time are likely to bias the model coefficients, although the magnitude and effect of any such bias is unclear without further analysis.

The passenger demand elasticity with respect to population is about 0.9, which implies that a 10 percent increase in regional population will lead to about a 9 percent increase in air traffic, other things being equal. *A priori* one would expect the demand elasticity for population to be around 1.0 or perhaps slightly higher. If the composition of the population is unchanged, then 10 percent more people would make 10 percent more trips. However, increasing population could imply changes in the composition of the population as well. People moving to the Portland area may be more likely to have family elsewhere and hence the need to make more air trips or to be working in those sectors of the economy that generate more business travel. This would tend to give a population elasticity higher than 1.0, but probably no higher than 1.05.

The per capita income elasticity is a little over 1.0 (1.04). This implies that air travel propensity will increase marginally faster than income. In fact, air passenger surveys show that air travel propensity increases very strongly with income. Households making $200,000 per year make significantly more air trips than those making $30,000 per year. Thus *a priori* one would expect air travel demand elasticity with respect to income to be higher than 1.04, although by how much would depend on any changes in the income distribution as well as the average.

If both population and income elasticity turn out to be higher in the future than predicted by the model, then the model will tend to under-predict the future growth in air travel for any given assumptions about the cost of air travel.

The elasticity with respect to PDX yield is –1.15 or slightly elastic (an increase in average air fares of 10 percent would reduce traffic by 11.5 percent). This is broadly consistent with airfare elasticity found in other studies, particularly for longer haul domestic trips.

A critical issue in the application of the model is the use of the dummy variable for years after 2001. This term reduces traffic by 32 percent below the level predicted by the model without including the variable. While it may be true that there are lingering effects of 9/11 that have continued to reduce the demand for air travel in recent years, the relevant question for the use of the model to forecast future traffic is how long these effects will last and whether the model provides a reasonable representation of the future causal relationships without the use of the dummy variable.

The recent trends in traffic suggest that keeping the dummy variable in the model for future years will tend to under-predict the level of air traffic for any given set of input assumptions. However, excluding the dummy variable completely is likely to greatly over-predict traffic levels, particularly in the near term, and anyway changing one term in a model will impact the values of the other coefficients.

- ❖ **Further work appears to be warranted to better understand the impact of the use of the dummy variables in the model on the forecasts, as well as the sensitivity of the forecasts to changes in demand elasticity with respect to population and income.**

Cargo Model

The air cargo model has been recently revised to respond to peer review comments on the previous model and thus has not been subject to the same level of review as the passenger model.

The revised cargo model is based on an underlying logistic growth relationship in the total weight of air cargo per thousand dollars of regional personal income (expressed in constant dollars) that the model development team has termed the cargo intensity. This relationship is adjusted by a term that reflects changes in the price of oil. In addition the model includes a dummy variable that reduces the predicted cargo tonnage by 28% for the years after 1999. The demand elasticity with respect to oil price is about –0.16, which appears reasonable from a consideration of the proportion of integrated air cargo carrier operating costs accounted for by fuel.

Figure 2 shows the fit of the model to the historical data.

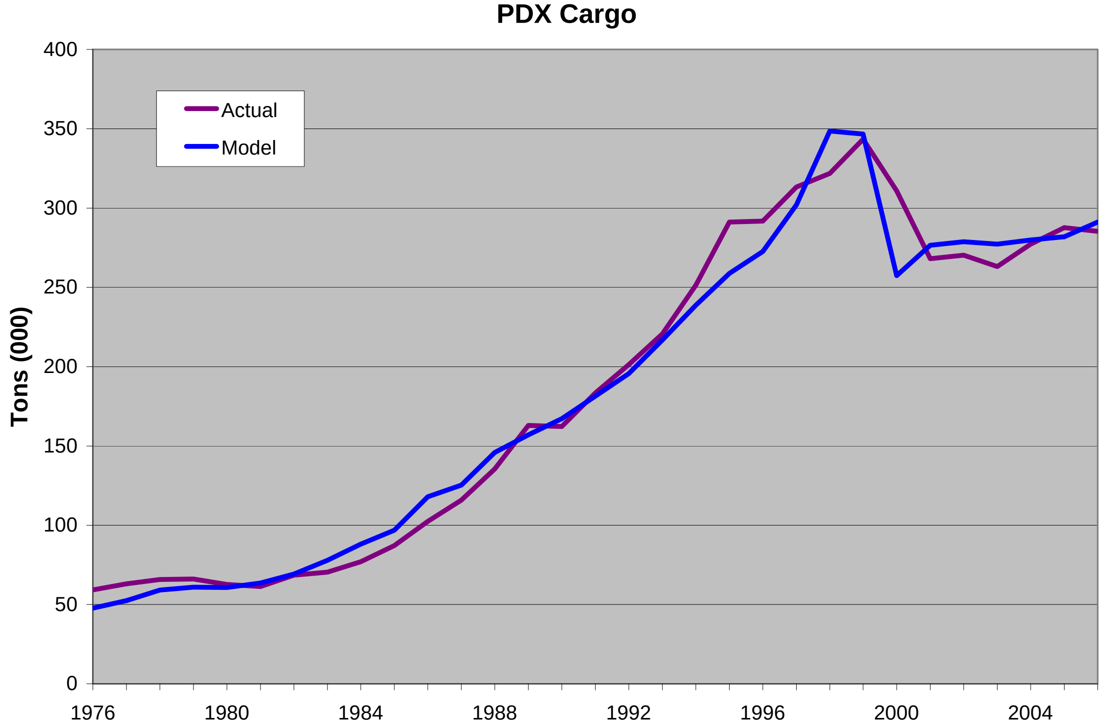

Figure 2 Total Air Cargo Model

It can be seen that the model fits the data fairly well until about 1993, but then does a relatively poorer job of explaining the year-to-year changes thereafter. The model captures the drop in traffic after 1999, but does so through a dummy variable. In fact the use of a single dummy variable for the period after 1999 overstates the drop from 1999 to 2000 and then predicts an increase from 2000 to 2001, instead of the decrease that actually occurred. Rather than predict air cargo tonnage directly, the model predicts air cargo intensity and then this is multiplied by the total regional personal income to predict the actual air cargo traffic. Not surprisingly, the corresponding fit of the cargo intensity model to the actual cargo intensity data looks very similar to the total air cargo results, as shown in Figure 3.

Figure 4 shows the underlying logistic growth relationship for the period of the model estimation and its projection to 2035. Since this relationship is defined solely in terms of the number of years since 1975, its future growth is unaffected by any causal variables, although for any given year, this relationship will be adjusted by the change in oil price.

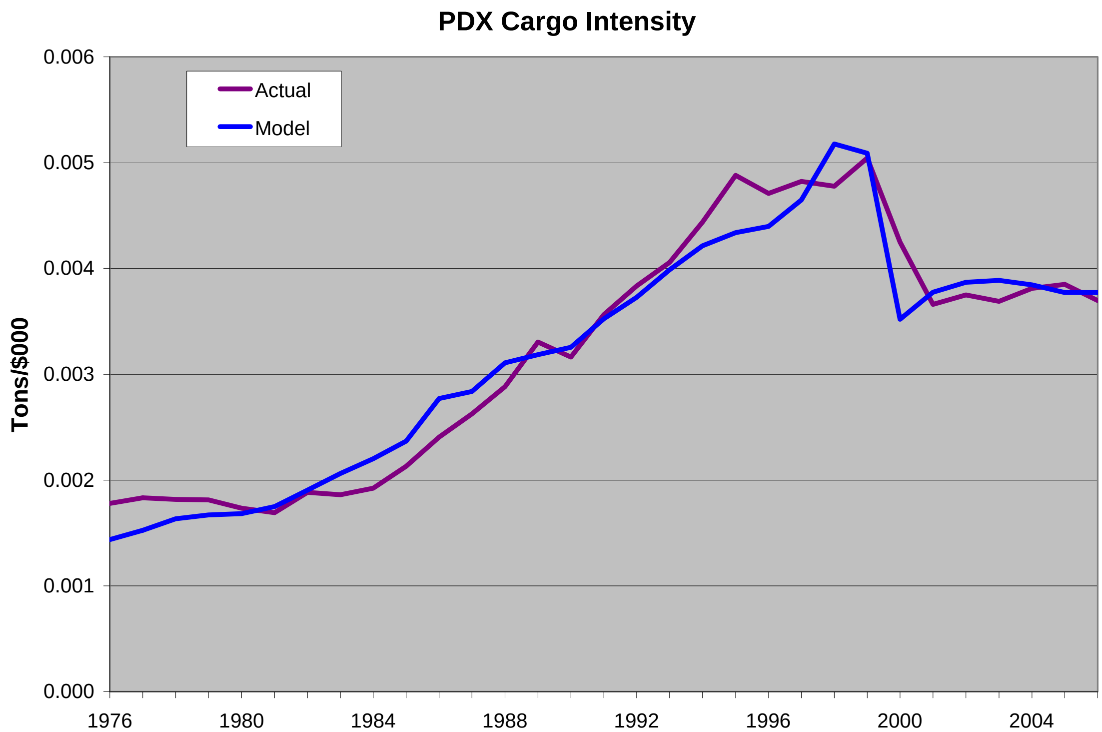

Figure 3 Air Cargo Intensity Model

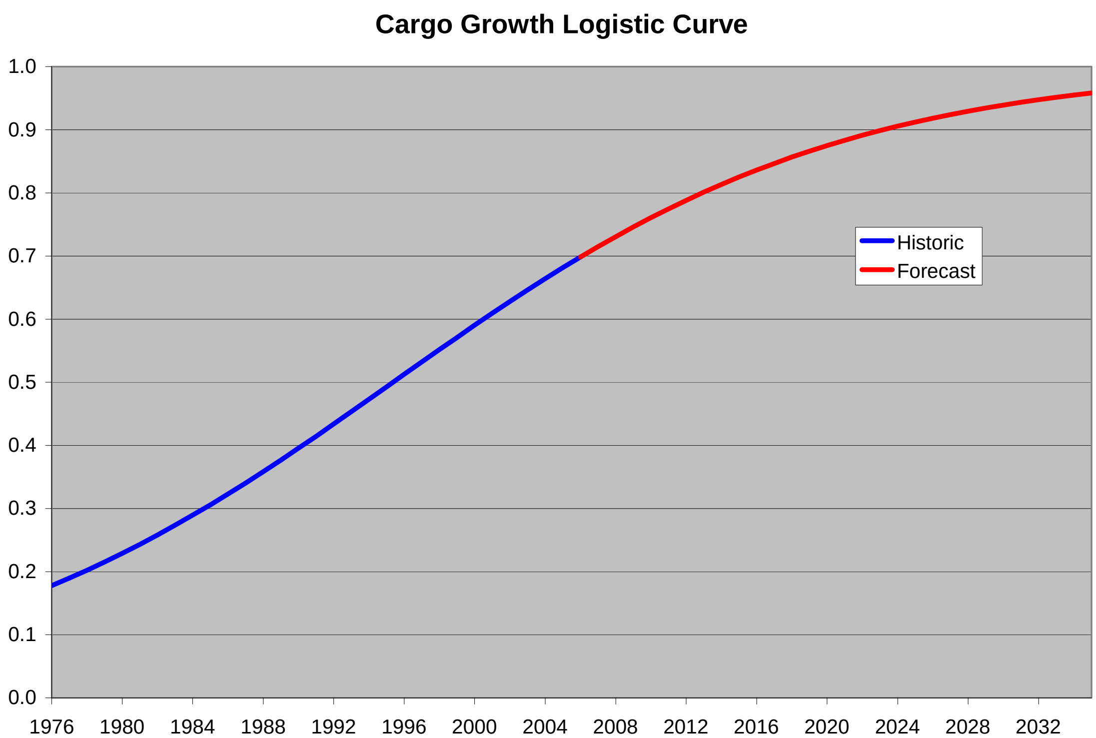

Figure 4 Logistic Growth Relationship

The forecasts of future air cargo traffic are significantly affected by the continued increase in the logistic growth relationship, which implies that the air cargo intensity will reach about 37% above 2006 levels by 2035. In deciding whether to accept the current forecasts, the PAG should consider whether this appears reasonable. The PAG should also decide whether the reduction in cargo intensity since 1999 is likely to continue in the future and by the same percentage.

- **Further work appears to be warranted to better understand the factors that caused the strong growth in cargo intensity until the late 1990s and the subsequent decline. This would allow a better informed assessment of the extent to which these relationships are likely to continue in the future.**

Development of Probabilistic Forecasts

A key feature of the forecasting approach adopted for the PDX Airport Futures project is the use of a probabilistic forecasting approach, also termed a risk analysis. This recognizes and attempts to address the inherent uncertainty in any attempt to predict the future. This uncertainty arises from at least three factors:

1. The models used in the forecasting do not completely explain the historical data on which they have been estimated, and as a consequence there is uncertainty in the values of their coefficients and residual errors in their predictions
2. Forecasts of future values of the explanatory variables are inherently uncertain, as illustrated by the difficulty of predicting future oil prices
3. Unanticipated events may occur that change the demand for air transportation (such as another terrorist attack) or structural changes may occur in the airline industry or the economy at large that change the relationships between the causal variables and air travel or air cargo traffic from those observed during the period on which the models have been estimated.

The way in which these factors are accounted for in the probabilistic forecasts is to define a probability distribution for each of the input variables in the models, and then calculate a large number of different values for each forecast traffic level, sampling randomly from these input variable distributions. The resulting values can then be used to define a probability distribution for the forecast traffic level. This process is termed a Monte Carlo simulation. In addition to varying the values of each input variable, the Monte Carlo simulation also varied the parameter values of the model and applied an error term to the model prediction, reflecting the uncertainty in the estimated model parameters and the residual errors in the model predictions of the historical data.

Obviously, this process is highly dependent on the assumptions made about the distributions of the future values of the input variables. These distributions cannot be directly observed from past data, but rather are measures of the confidence with which we can predict the future values of the input variable. The assumptions made in developing the distributions for use in the Monte Carlo simulation are discussed further below in the section on forecast assumptions.

An important consideration is performing a Monte Carlo simulation is whether the distribution of each input variable is assumed to be independent or whether the future values of some of these variables are likely to be correlated. For example, a higher oil price may be more likely to be associated with a low growth rate in per capita income than a high growth rate. In the current analysis it was assumed that each variable is independent. With the current models this will tend to produce a smaller spread of values than if the variables were assumed to have some correlation.

Of the explanatory variables included in the forecast models, the one that has the highest level of uncertainty about future values is the average value of airline yield (in the case of the passenger model) or the price of oil in the case of the cargo model. Since the price of oil has a direct effect on the passenger yields needed for the airlines to remain economically viable, the analysis developed a sub-model to predict the effect of changing oil price on airline yield. This considered not only the effect of future oil prices on airline fuel costs, but also likely changes in non-fuel costs and airline load factors, since these also affect the yields needed for the airlines to cover their operating costs.

One factor that will affect future airline fuel costs that was not explicitly included in the analysis is the continuing trend in improved aircraft fuel efficiency. The analysis implicitly assumed that this is subsumed in the past relationship between oil prices and airline fuel costs. However, any such relationship is likely to be very weak, given past trends in aircraft fuel efficiency and oil prices.

- ❖ **Further work appears to be warranted to clarify how a more explicit accounting for trends in aircraft fuel efficiency might change future projections of airline yield.**

Derivation of Forecasts of Aircraft Operations

The forecast models predict total enplaned passengers and total air cargo tonnage. In order to translate these forecasts into forecasts of aircraft operations, a process was developed to allocate the forecast traffic to individual markets and then estimate the number of aircraft operations needed to serve each market. This process considered the airlines currently serving each market, the aircraft fleet that they currently operate and any anticipated future changes in their fleet, and trends in average load factors.

This process was based on a number of key assumptions that should be borne in mind in interpreting the aircraft operations forecasts:

- The volume of international passenger travel was forecast externally from the forecast model by making assumptions about future growth rates in international travel under a median, low and high growth scenario. The future level of domestic travel was then obtained by simply subtracting the international travel from the total travel predicted by the forecast model.
- The share of passenger travel in each market served by non-stop flights from PDX was assumed to remain the same as in 2007.
- The share of the passenger traffic in each market carried by each airline was assumed to be the same as in 2007.

- The future percent change in load factors in each forecast year was assumed to be the same for every airline and market, although the growth in load factor compared to 2007 was assumed to increase slowly over time.

These assumptions implicitly assume that the basic structure of the airline industry remains unchanged from 2007 and in fact is independent of changes in the determinants of the future demand for air travel, particularly the price of oil. While it is true that airlines cannot change their fleets significantly in the short run, they can certainly do so over the time frame of the forecast. It can be expected that continuing high oil prices will force airlines to use larger, more fuel-efficient equipment and may lead to shifts in airline market share in favor of those airlines with more fuel-efficient aircraft. Pressures to increase average aircraft size may also result from airport capacity concerns under scenarios of high traffic growth, as illustrated by the current discussion about the introduction of congestion pricing at capacity-constrained airports.

While accounting for these effects at the market level may be beyond the resources of the current project, the limitations that this imposes on the resulting forecasts of aircraft operations should be understood.

Forecast Assumptions

Forecast values of future air traffic are only as good as the assumptions on which they are based. In the case of the current forecasting process, there are two separate issues that should be considered: the assumed expected value for each input variable for a given forecast year and the assumed distribution of likely values around the expected value. In fact, the expected (or average) value is implied by the distribution. However, it is easier to think about the assumptions by separating the two issues.

Expected Values of Explanatory Variables

There are really only three explanatory variables that drive the PDX air traffic forecasts: population, per capita income, and the price of oil. Although the passenger model is based on the average PDX yield, the projected values of this variable were derived from the projected price of oil.

The expected future values of population and per capita income were obtained from regional forecasts prepared by Portland Metro. These forecasts appear to be widely accepted and there is obviously a good argument to be made that planning for PDX should be based on consistent assumptions with the regional planning for the rest of the transportation system. On the other hand, the Port may wish to understand the implications for its own financial and facilities planning of different assumptions about future growth in population and income. While these assumptions are varied in the probabilistic forecasts, it is not currently possible to separate out the effects of varying one particular assumption. This would require a sensitivity analysis with different input assumptions.

The expected future values for oil price were based on the most recent U.S. Department of Energy projections, although an extensive comparative analysis was undertaken of oil price forecasts by other organizations. In recent years, the U.S. Department of Energy has significantly under-predicted the increase in the real price of oil. Whether this tendency to under-predict the future price of oil will apply to the most recent forecast only time will tell.

However, this places a particular importance on the assumed distribution of likely future values, as discussed further below. A sensitivity test of oil price assumptions indicated that a 21% increase in the price of oil resulted in about a 4% decrease in passenger traffic, which seems reasonable given the proportion of airline costs accounted for by fuel and the elasticity of demand with respect to yield. This suggests that a doubling of oil price over the DOE forecasts could reduce traffic by perhaps 20%. However, the more serious concern from the perspective of the forecasts is that if the lower-bound scenario is too low ($44 per barrel in 2035 in 2006 dollars), this will tend to overstate future growth in traffic.

Other assumptions that influence the forecast include future values of average load factor and aircraft size. These were generally based on the most recent national aerospace forecast prepared by the Federal Aviation Administration (FAA). These values are not likely to deviate significantly at a national level from the FAA projections, since there are physical limits on how many seats can be sold or how quickly the airlines can increase the average size of their fleets. Although the corresponding values at PDX may differ somewhat from national averages, due to differences in traffic composition, airline shares and aircraft fleets, the growth trend is likely to be similar, since the airlines and flights serving PDX are part of the larger system. In any event, differences in these factors are likely to have much less impact on the forecasts than changes in the price of oil.

Impact of Climate Change Policies

One issue that has been of particular interest to the Planning Advisory Group Forecast Subcommittee is the effect of future climate change policies on the cost of air travel or the propensity to make air trips. These effects have been incorporated in the forecasts through adjustments to the assumed price of oil. These adjustments have been expressed as a carbon tax that would be payable by oil companies (and presumably passed on to their customers) or directly by oil users. Future costs of carbon emissions could also be the result of a cap-and-trade program or other regulatory environment that requires airlines to purchase carbon offsets. While the need to purchase carbon offsets will not affect the price of oil directly, to the extent that the required offsets will be based on fuel consumption, the effect will be the same as if the price of oil increased.

In addition to assuming a price per barrel for a carbon tax, the analysis also assumed that only a proportion of this tax would be passed on the airlines. This obviously would not apply to a carbon offset program, unless the government were to provide exemptions to airlines on some basis. It is also unclear why oil companies would choose to pay some proportion of a carbon tax out of their profits (or indeed could even afford to). It is possible that the introduction of a carbon tax would reduce the demand for oil, resulting in oil producers lowering their prices to offset the lost revenue through greater production. However, this involves a complex market dynamic that is beyond the scope of the current forecast analysis to perform. It is also unclear to what effect these factors have already been accounted for in the DOE scenarios. Of course, as far as the traffic forecast is concerned, there is no difference between a higher carbon tax, only some of which is passed on to the airlines, and a lower tax that results in the same cost per barrel. On balance, it would be clearer to assume that the entire carbon tax is passed on the airlines and simply vary the assumed tax.

The assumed price of carbon emissions ranged from $35 to $109 per metric ton in 2035 in 2006 dollars, based on a recent study by the Massachusetts Institute of Technology. However, there are those who feel that much higher levels will be required to achieve the targets proposed in recent state and federal legislation, particularly by 2035 when many of the cheaper alternative strategies to reduce carbon emissions will have been exhausted. In addition, the assumed carbon prices do not appear to account for the radiative forcing effect of aircraft emissions, which has been estimated to require the direct aircraft emissions (expressed as carbon equivalent) to be increased by a factor of about 2.7 or more to fully offset the impact on climate change. In addition, there is no adjustment in the models to reflect possible effects of growing public concern about climate change on the propensity to make air trips. **Thus it appears that the assumptions used in the forecast analysis may significantly understate the likely future costs of climate change policies or public attitudes.**

- **Further work appears to be warranted to clarify the sensitivity of the forecasts to more stringent assumptions about potential responses to climate change.**

Assumed Distributions

The assumed distribution for future values of the regional population was obtained directly from the population forecast approach used by Metro. However, in the case of the per capita income Metro had not developed a predicted distribution and it was necessary to develop an assumed distribution from the per capita income forecast scenarios prepared by Metro. This distribution was based on analysis of past performance of the Metro income forecasts by the Jacobs Consultancy team. Since this found that Metro had tended to over-predict the growth in real income in the past, this was incorporated in the assumed distribution. A consequence of this approach is that the PDX traffic forecasts are based on a somewhat lower average per capita income assumption than the baseline forecast by Metro.

The assumed distribution of future values of the price of oil was based on the high and low scenarios forecast by the U.S. Department of Energy (DOE). No adjustments were made to account for the past performance of the DOE in forecasting oil prices. Therefore if the DOE high growth scenario has substantially under-estimated the potential for continued future growth in oil prices, this possibility will not be considered in the PDX traffic forecasts. Thus while the assumed distribution for the price of oil covers a wide range (from $44 per barrel to $128 per barrel in 2035 in 2006 dollars), this range may understate the possibility of even higher prices.

An important assumption in developing distributions of likely future values of explanatory variables is the shape of the distribution. Apart from the distribution for population, which was obtained from Metro, the other distributions were assumed to be triangular. In most cases, the upper and lower limits of the distribution were derived from high and low scenarios of the relevant variable. This constrains the values to lie between the high and low scenario. It is likely that in developing high and low growth scenarios, other organizations are not assuming that these are the highest and lowest values that could possibly occur, but rather represent a range within which future values have a reasonable likelihood of occurring (although what this likelihood is may not be explicitly defined). **This will tend to under-predict the range of forecast values, although it is unclear by how much.**

- **Further work appears to be warranted to clarify the sensitivity of the forecasts to changes in the assumed distributions for the explanatory variables.**

Use of Probabilistic Forecasts in Master Planning

The development of probabilistic forecasts that express the likelihood of future traffic levels exceeding any given value raises the question of how such forecasts can be used in the remaining steps of the master planning process. There are two important points that should be noted in discussing the probabilistic forecasts:

1. The forecasts express the likelihood of obtaining a future traffic level that lies between any two values, or conversely the likelihood of the future traffic being greater than (or less than) a given value. It is misleading to talk about the “most likely” value for the future traffic level in any given year. The median value (50 percentile) represents the value that the future traffic is equally likely to be greater than or less than. While the likelihood of the future traffic lying in a given range around this value (say plus or minus 1 million annual passengers) may well be higher than the corresponding likelihood of the traffic lying within the same range around any other value, this likelihood is still very small unless the range in question is fairly wide.

2. Although the graphs of forecast traffic levels show curves for different percentile forecasts, this does not imply that the future traffic growth will follow any one of these curves. These merely join points with the same percentile likelihood in successive years. It is entirely possible for the future traffic to evolve such that it corresponds to one particular percentile value in one year and as quite different percentile value some years later. Obviously in the real world there are constraints on how rapidly traffic can grow, so a change from the 10-percentile value to the 90-percentile value in just a few years is highly unlikely.

The real value of the probabilistic forecast, apart from being more transparent about how uncertain the forecasts are, lies in the ability to address questions such as “how soon might we need the capability to handle 30 million annual passengers at PDX?” or “what is the likelihood that we will need to build another terminal concourse by 2015?” While some alternative facility development options may be fairly easy to advance or defer if traffic grows more quickly or more slowly than assumed, others may be much harder to implement sooner than expected due to the need for other actions to be completed that take time (e.g. land acquisition or relocating other facilities).

One useful way to interpret the graphs that show the varying percentiles of forecast traffic level for any given year is to take a horizontal line at any given traffic level and express the percentiles in terms of the corresponding likely future year at which the given traffic level will be reached.

Considerations for Further Work

The various issues identified in this review suggest that a number of the steps in the forecast process could benefit from further study and development, and that further examination of some of the assumptions appears warranted.

While this additional work need not delay starting the remaining steps in the master planning process on the basis of the current forecasts, it would be in the Port’s and City’s interests to continue to pursue these issues in parallel with this other work, so that there is a better understanding of the sensitivity of the current forecasts to the issues identified in the review as the other steps in the master plan evolve. It would also be desirable to continue to develop the forecasting methodology so that future updates of the forecasts can benefit from better models and analytical tools. This will be particularly helpful when the forecasts are revisited closer to the time when important facility development decisions need to be made. However, deferring further work to that time will not only leave these issues unresolved for several years, but runs the risk of not leaving enough time to complete the necessary analysis before a decision is needed. Where additional data is required, as in the case of air cargo, it is important to begin collecting this as soon as possible.

The more significant aspects that could benefit from further work include:

- Improving the passenger and cargo forecast models to better account for recent changes in the airline industry and reduce the importance of dummy variables in predicting recent traffic trends. These improvements should include the development of a market-based analytical approach so that future changes at a market level can be explicitly forecast and do not need to be assumed.
- Incorporation of a more detailed representation of airline costs that can better account for the effects of fuel prices, aircraft fuel efficiency, and changes in other non-fuel costs.
- Obtaining more detailed regional cargo data to identify changes in the relative role of trucks and air, and trends in air cargo use by different sectors of the economy.
- Further study of the likely impact of climate change policies on oil prices or airline costs, as well as possible effects of growing public concern about climate change on the propensity to make air trips.
- Further study of the most appropriate probability distributions for the forecast values of explanatory variables in the Monte Carlo simulation and consideration of likely correlation between future values of explanatory variable.
- Development of analytical software to simplify the process of re-running the models and forecast analysis with different assumptions. The current analysis process is very labor intensive and relies heavily on the use of spreadsheet programs, which are not readily usable by others.

While this work would require a significant commitment of resources, this would still be fairly minor compared to the future capital investment in airport facilities, and indeed compared to the engineering and design work for those facilities. It would also greatly enhance the Port's ability to track the changes in airport traffic compared to the forecasts and understand whether the level or pattern of demand is changing in unforeseen ways that should precipitate a review of the implications for its established master plan.