Estimating the Cost of Commercial Airlines Catastrophes—A Stochastic Simulation Approach

Romel Salam, FCAS, MAAA

Estimating the Cost of Commercial Airlines Catastrophes - A Stochastic Simulation Approach

Romel Salam, FCAS, MAAA

Abstract

Actuaries are increasingly finding more applications for stochastic simulation in pricing, reserving, DFA, and other insurance and financial engineering problems. For instance, stochastic simulation has gained acceptance as a pricing tool for property catastrophe coverage in the insurance, reinsurance, broker, and investment communities. This has required primary companies to compile and provide information at a more detailed level than they did only a few years ago. Various commercial simulation products have emerged to help companies assess and price their property catastrophe exposures. Although there are many parallels between the catastrophe exposures of property and commercial aviation risks, the use of simulation is not widespread in the assessment of commercial aviation catastrophic exposures. In this paper, we present the framework for a simulation model for commercial aviation catastrophes and we discuss various aspects of designing such a model including the level and type of information needed.

Acknowledgement

We would like to thank Yipei Shen and Pavel Zhardetskiy for their many contributions to this project. We are also indebted to many of our other colleagues at Transatlantic Reinsurance Company who, directly or indirectly, knowingly or unsuspectingly, willingly or under physical threat, contributed to this paper. Of course, we remain solely liable for mistakes, misrepresentations, and any libelous material that appear in the following pages.

Introduction - The need for a stochastic model

The claims covered by a comprehensive commercial airline policy can be broken into two groups. The first group consists of trivial claims such as lost luggage, "slip and fall" accidents, or minor damage to the hull of an aircraft while the second group comprises catastrophic claims arising out of airplane crashes resulting in serious injuries, fatalities, property damage, and major or total loss of an aircraft. Most of the pricing tools that are used to price airline's hull and liability exposures tend to rely on experience rating techniques. Under a basic experience rating method, the projected losses are based on an average of past losses adjusted for trend and development. An experience rating approach may work relatively well when only the non-catastrophic exposure of airlines is considered. However, traditional experience rating methods would tend to overstate the expected loss when one or more catastrophes are included in the experience period, and, conversely would tend to understate the expected loss when there are no catastrophes in the experience period. Under a more sophisticated experience rating approach, losses are separated into their catastrophic and non-catastrophic components. The catastrophe losses are then compiled and averaged over a very long period of time in order to come up with an "expected catastrophe loss amount" similar to what is used in property ratemaking. Even under the latter approach, the question needs to be asked as to whether past catastrophe experience is representative of future experience. First, the frequency of catastrophic accidents may have changed over time due to such factors as improved aviation technology, better or worse safety regulation, or increased air traffic. Secondly, the costs of hull and liability coverage are indeed impacted by not only general inflationary trends which can be reflected within a traditional experience rating model, but also by changes in an airline's fleet, passenger load factors[1], destination and passenger profiles[2] which are harder to reflect in an experience rating exercise. Finally, as we look to the future, the introduction of new aircraft models such as the Airbus 380 model - which could

[1] Passenger load factor represents the average percentage of an airline's seating capacity that is filled.

[2] Destination profile for an airline refers to the countries to which the airline is flying. The liability damage award of accident victims may vary by country. For instance, a priori, an airline operating domestic flights solely in India would have a lower liability potential than one operating solely in the United States. Passenger profile refers to the age, occupation, income of passengers as these are all factors that can determine the level of compensatory damage of accident victims. A priori, an airline whose core clientele was made up of college freshmen going on vacation would have a lower liability potential than one whose core clientele consisted of well paid corporate managers going to business meetings.

accommodate up to 840 passengers in a single class configuration - to an airline's fleet, the addition of new set of routes and destinations, the continued evolution of contracts and laws establishing the compensation of victims of airline accidents will all make it less likely that traditional experience rating will remain an adequate forecasting tool.

The stochastic model that we present avoids most, if not all, of the pitfalls of traditional experience rating methods. While judiciously making use of historical data such as past accident rates, the model will rely on the most current information relating to an airline's fleet, passenger and destination profiles, number of departures (or miles flown), and passenger load factors. The model will also be flexible enough to allow the user to incorporate his/her views on the impact of legal changes on the cost of liability, hull or other related costs of accidents. This model will be especially well suited for analyzing contracts which carry a lot of bells and whistles. A simulation model that breaks down the loss process into its many components also forces the user to think about the different factors that impact on the costs of airline catastrophes. Perhaps, one drawback of such a model is that it requires a more detailed level of information than generally needed in performing an experience rating exercise. However, such information is generally available with a little bit of research.

This paper is organized into eight sections. In section 1, we present a schematic of a stochastic simulation model for evaluating the cost of passenger liability coverage. In section 2, we define commercial airlines and airline catastrophes. In section 3, we delve into the area of frequency, including the choice of an appropriate measure of exposure. In that section, we also explore the issue of classification by relying on work presented in "Reinventing Risk Classification – A Set Theory Approach" [6]. We revisit a statistic introduced in that paper for the purpose of making inferences about Poisson distributed events. We use this statistic to comment on a Wall Street Journal article, which sought to discuss the relative safety of several aircraft models. We then tackle the issue of whether the rate of airline catastrophes has changed at all over time in the same section. In section 4, we look at the cost of catastrophes for different coverages, including more easily determined costs such as those for hull coverage to more challenging ones such as passenger liability, third party liability, and products liability. We also briefly touch on the issue of classification relating to passenger liability costs in that section. In section 5, we discuss how various results coming from the model can be validated against actual historical data. In section 6, we offer some thoughts on how to incorporate the risks of terrorism or sabotage in the model. In section 7, we give an example of the simulation model for a cover for a hypothetical group of airlines. In section 8, we provide final thoughts.

1) Simulation Scheme

Figure 1.1 below shows how we would generate passenger liability losses using a simulation model. This scheme would vary depending on the level of information available and the coverage of interest.

2) Definitions

Before we go too far into this discussion, let's try to agree on the topic of discussion itself by attempting to put some parameters around two of the terms that are central to this paper:

2.1) Commercial Airlines

Insurance underwriters generally differentiate between commercial and general aviation. General aviation typically encompasses operation of smaller airplanes used for leisure, industrial and agricultural purposes, or simply in the private transportation of individuals or employees. Helicopter and balloon operations are generally lumped into the general aviation category. Commercial aviation involves the transportation, for compensation or hire, of persons or cargo by aircraft. In the US, a commercial operator is one that has been certificated by the Federal Aviation Administration (FAA) under Code of Federal Regulation (CFR) part 121 (airlines) or CFR Part 135 (commuters) to provide air transport of passengers or cargo. So-called air taxis and commuters operate smaller aircrafts and carry few passengers per flight whereas airlines typically operate jet aircrafts that can carry large loads of passengers per flight. More recently, the line between commuters and airlines has been blurred by acquisitions as well as the amendment of some of the FAA codes. This paper is concerned mostly

Figure 1.1 – A simulation framework

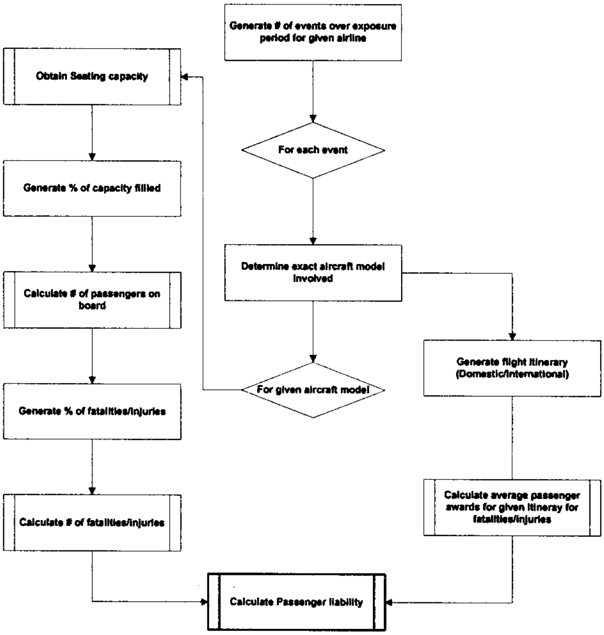

with commercial airlines, as most of the statistics we will discuss will relate to airlines represented by US operators certificated under CFR Part 121 and operators in other countries with similar certification.

2.2) Airline Catastrophes

We have referred several times already to airline catastrophes as if the term was self-explanatory. In fact, one might take several views as to what constitutes a catastrophe. One perspective of catastrophes could be that of an excess of loss reinsurer who would typically be impacted only by events above a certain threshold. For instance, a reinsurer could *define a catastrophe as an accident, occurring between takeoff and landing, involving one or several aircrafts, and which results in major damage to or destruction of an aircraft's hull.* Under this definition, for instance, damage to an aircraft from a hailstorm or an earthquake while garaged would not be counted as a catastrophe, neither would fatalities or injuries occurring as a result of air turbulence, food poisoning, or falling luggage. A midair collision between two or more aircrafts would be counted as one catastrophe. Throughout this paper, we will use slightly different definitions of catastrophe and different data sources to illustrate different aspects of the simulation model. The exact definition used is of no particular importance since we are not trying to promote any one definition but rather trying to present a method by which the cost of such catastrophes, however defined, can be evaluated. It is, however, important that the data collected for the purpose of constructing a model be consistent with the definitions used in the contracts and products that are being evaluated.

3) Modeling the Frequency of Airline Catastrophes

How do we best model the number of airline catastrophes? Within the casualty and property actuarial practice, there are two distributions that are commonly used to represent the frequency distribution of accidents, namely the Poisson and the Negative Binomial distributions. We, a priori, will work with the Poisson distribution because of its simplicity and its intuitive appeal[3]. A modeler is free to choose other distributions that might work better or as well.

[3] There are three postulates implied by a Poisson process:

The Poisson model is defined as follows: $f(x) = \frac{(\lambda d)^x e^{-(\lambda d)}}{x!}$,

where x is the number of catastrophic accidents, λ is the expected rate of accident per exposure unit, and d is the number of exposure units for the period under review.

3.1) Picking an exposure base

The exposure base as will be used here is the unit upon which frequency will be measured. Airlines usually report the number of departures, miles (or kilometers) flown, hours flown for annual, quarterly, and even monthly periods. Any of these measures could serve as an exposure base since they are almost perfectly correlated. The modeler's decision as to which of these potential units of exposure to use may be based on which is found to be more accessible, more accurate, measured and defined more consistently overtime. The modeler needs to be well aware of potential distortions in the exposure data especially when using different sources to gather the data. We will use the number of departures as our unit of exposure because of some evidence showing that the risk of catastrophic occurrences is concentrated around takeoffs and landings. We, however, have not found any significant differences in our results when we used hours or miles flown as measures of exposure.

3.2) Classification

For the twenty-two year period from 1980 to 2001, commercial airlines catastrophes[4] occurred in the world at a frequency of 1.22 per million departures. Should we look at this frequency rate as being applicable to all commercial airlines and use it as the basis for the λ in the Poisson model for all airlines? This approach could potentially result in the underestimation of the accident risk for some groups of airlines while resulting in the overestimation of that risk for some other groups.

---

1) The numbers of occurrences in non-overlapping time intervals are independent.
2) The number of occurrences in a time interval has the same probability for all intervals.
3) The probability of two or more events in a small time interval is zero.

We believe that the occurrence of catastrophic airline accidents, excluding those caused by willful acts, satisfy all three postulates.

[4] Catastrophe here is defined as accidents resulting in total destruction of an aircraft. Data comes from a proprietary source and is based on Western-built aircrafts only.

Conversely, should we group airlines into cells[5] along certain risk characteristics[6] and proceed to calculate frequency rates for each cell based on the data within? Are we to then presume that airlines across these cells have fundamentally different rates of catastrophic accidents by virtue of our having devised the grouping scheme? Under this approach, we risk assigning different, perhaps even significantly different, frequency rates to cells of airlines that have essentially the same propensity for catastrophic accidents. Let's assume for a moment that airlines are grouped based on the subcontinent on which they are domiciled. In the twenty-two year period from 1980 to 2001, the average frequency of catastrophic accidents for North-American airlines has been around 0.5 per million of departures while that for Western-European airlines has been closer to 0.6 per million of departures. Did the difference in the observed accident rates arise out of the random nature of catastrophic accidents or did it arise out of a fundamental difference in the propensity of accident for the two groups? The search for answers to these questions spun a classification methodology introduced in the paper titled "Reinventing Risk Classification – A Set Theory Approach" [6]. We will use the procedures from this paper to give an example of a classification scheme for airlines. Before we do, however, we want to reacquaint the reader with a statistic, $\hat{R}_0$, introduced in the aforementioned paper and which was used to make inferences about Poisson distributed populations.

3.2.1) A review of the $\hat{R}_0$ statistic.

Let λ_A and λ_B represent the expected frequency rates for two Poisson populations A and B, with d_A and d_B units of exposure, respectively. Also, let $\hat{\lambda}_A$ and $\hat{\lambda}_B$, represent the maximum likelihood estimates of λ_A and λ_B, respectively. In [6, p. 105-114], we show that if $\lambda_A = \lambda_B$,

$$\hat{R}_0 = \frac{\hat{\lambda}_A - \hat{\lambda}_B}{\sqrt{\frac{\hat{\lambda}_A}{d_A} + \frac{\hat{\lambda}_B}{d_B}}} \rightarrow N(0,1) \text{ for large}^7\ d_A \text{ and } d_B \text{ values.}$$

[5] A cell is a set of airlines with the same risk characteristics [6, p. 89].
[6] Risk characteristic is an attribute that identifies a risk or group of risks [6, p. 88].
[7] i.e. as d_A *and* $d_B \rightarrow \infty$ [6, p. 105-114].

We then use $\hat{R}_0$ to make inferences about the equality of the frequency rates underlying pairs of Poisson distributed populations. If we define the null hypothesis as $\lambda_A = \lambda_B$, then we will reject that hypothesis at the 10% significance level (90% confidence interval) if $\hat{R}_0$ falls outside the range of (-1.65, 1.65). We explain in [6, p 94] that $\hat{R}_0$ can be thought of as the observed distance between the two populations' samples. If that distance is small, we tend to accept the hypothesis that the populations have the same expected frequency. If it is large, we tend to reject the equality hypothesis. Observe that $\hat{R}_0$ depends not only on the MLE estimates of the population frequencies but also on the number of exposure units of the respective populations. For instance, the absolute value of $\hat{R}_0$ increases as the number of exposure units increases (everything else being equal).

A Wall Street Journal article [5] in which the author sought to demonstrate the poor safety record of the MD-11 aircraft relative to other similar models provides a good example of how $\hat{R}_0$ can be used to make inferences about Poisson populations. The article shows a graphic with accident rates by airplane types, which we summarize in table 3.1 below.

If we assume that the number of accidents for each of the aircraft models is Poisson distributed, we can use $\hat{R}_0$ to make inferences about the relative safety of these models. The exposure units are the number of millions of departures, while the maximum likelihood estimates of the expected frequencies are represented by the frequency per million departures. For instance, to test the hypothesis that the underlying accident rates for the MD-11 and the A300-Early are the same, we calculate $\hat{R}_0$, using an algebraic equivalent[8] of the formula introduced in section 3.2.1, as follows: $\hat{R}_0 = \frac{6.54 - 1.29}{\sqrt{\frac{6.54^2}{5} + \frac{1.29^2}{7}}} = 1.771$.

Table 3.1- Accident Rates by Airplane Type[9]

Aircraft Model	(1)=(2)/(3) Million of Departures	(2) Hull Losses[10]	(3) Frequency per million Departures
B-707/720	17.8	115	6.46
DC-8	12.2	71	5.84
B-727	72.2	70	0.97
B-737-1 & 2	50.4	62	1.23
DC-9	58.1	75	1.29
BAC 1-11	8.3	22	2.64
F-28	8.1	32	3.94
B747-Early	11.1	21	1.90
DC-10	7.8	20	2.57
A300-Early	5.4	7	1.29
L-1011	5.2	4	0.77
MD-80/90	23.3	10	0.43
B-767	7.3	3	0.41
B-757	8.7	4	0.46
Bae146	5.1	3	0.59
A-310	2.9	4	1.40
A-300/600	2.2	3	1.34
B-737-3, 4 & 5	30.8	12	0.39
A-320/319/321	7.3	7	0.96
F-100	3.8	3	0.80
B747-400	2.0	1	0.49
MD-11	0.8	5	6.54

[8] $\hat{R}_0 = \frac{\hat{\lambda}_A - \hat{\lambda}_B}{\sqrt{\frac{\hat{\lambda}_A}{d_A} + \frac{\hat{\lambda}_B}{d_B}}} = \frac{\hat{\lambda}_A - \hat{\lambda}_B}{\sqrt{\frac{\hat{\lambda}_A^2}{n_A} + \frac{\hat{\lambda}_B^2}{n_B}}}$ where n_A and n_B represent the number of accidents for populations A and B, respectively, and $\hat{\lambda}_A = \frac{n_A}{d_A}$ and $\hat{\lambda}_B = \frac{n_B}{d_B}$.

[9] The article lists Boeing as the source of the data. The exposure units (million of departures) were not provided but calculated as the ratio of the number of accidents to the accident rate.

[10] The article defines hull losses as damage so severe the plane isn't repaired.

At a 10% significance level, we would reject the hypothesis that the two aircraft models have the same propensity for accident since $\hat{R}_0$ falls outside the interval (-1.65, 1.65). However, we would not be able to reject the hypothesis at a 5% significance level since $\hat{R}_0$ falls in the interval (-1.96, 1.96). In table 3.2, we calculate the $\hat{R}_0$ values between the MD-11 model and other aircraft models in table 3.1.

Table 3.2 - $\hat{R}_0$ values between $MD-11$ and other models

Aircraft Model	$\hat{R}_0/MD-11$
B-707/720	0.027
DC-8	0.233
B-727	1.903
B-737-1 & 2	1.813
DC-9	1.793
BAC 1-11	1.309
F-28	0.865
B747-Early	1.571
DC-10	1.332
A300-Early	1.771
L-1011	1.956
MD-80/90	2.087
B-767	2.089
B-757	2.072
BAe146	2.021
A-310	1.709
A-300/600	1.719
B-737-3, 4 & 5	2.101
A-320/319/321	1.893
F-100	1.939
B747-400	2.040

Before drawing any conclusions from the above table however, one would need to look into other factors that may impact on the accident rates. For instance, if the MD-11 losses were coming disproportionately from a particular operator or group of operators, the issue might be more specific to the operator or group of operators rather than to the aircraft model itself.

3.2.2) A classification scheme for the frequency of airline catastrophic accidents

A priori, one might expect airlines operating under similar jurisdictions and having similar types of operations, fleet, staff training, and safety procedures to display the same expected rate of accidents. The jurisdictions – which could be tantamount to countries – help explain the degree of oversight to which airlines are subject, the adequacy and competence of air traffic control, the level of competition in the market, the resources available to regulators to enforce safety rules, the degree of accountability of regulators and airlines to the public, and the public's attitude toward safety. Factors that may help delineate among jurisdictions include political system, economic standing, and judicial/tort system. Factors that may explain differences between airlines within the same jurisdictional group include size and years of operation. For instance, to the extent that there are economies of scale present in aircraft maintenance or staff training, larger airlines may exhibit a better safety record than smaller ones. For illustration purposes only, let's look at a two-dimensional classification scheme where jurisdiction and size of operations are the two classification variables. Then, we will define five jurisdictional groups and three sizes, which will result in fifteen cells for which the exposures and MLE estimates are shown in tables 3.3 and 3.4 below:

Table 3.3 – Departures in millions

	Jurisdiction 1	Jurisdiction 2	Jurisdiction 3	Jurisdiction 4	Jurisdiction 5
Large	56.9	120.8	-	4.8	15.8
Medium	33.5	8.9	7.4	12.5	15.5
Small	3.5	2.2	2.8	2.6	2.0

Table 3.4 – Initial MLE Estimates (Accidents / Million Departures)

	Jurisdiction 1	Jurisdiction 2	Jurisdiction 3	Jurisdiction 4	Jurisdiction 5
Large	0.527	0.356	N/A	3.120	2.019
Medium	0.507	1.679	4.305	2.800	2.710
Small	1.443	2.736	17.852	9.237	4.085

A revised set of estimates is obtained in table 3.7 below based on a procedure introduced in "Risk Classification – A Set Theory Approach" [6] and at the presentation of the paper at the winter 2002 meeting. Basically, each cell[5] defines a class[11] made up of the cell itself and possibly of other cells that

[11] Please refer to [6, p. 88 -90] for a definition of these terms.

are compatible[11] with it. Once the classes are defined, each cell within a class is given a credibility[11] weight and the revised estimate for each cell is the credibility weighted average of all the cells in its class. The classes defined by each cell and the credibility weights assigned to all the cells in each class are shown in tables 3.5 and 3.6 below, and again in exhibits 5 and 6 of Appendix A. The steps leading to table 3.5, 3.6, and 3.7 are detailed in Appendix A.

Table 3.5 - Classes defined by each cell

Cells	Classes
J1/L	{J1/L, J1/M, J1/S}
J2/L	{J2/L}
J4/L	{J4/L, J4/M}
J5/L	{J5/L}
J1/M	{J1/M, J1/L, J1/S}
J2/M	{J2/M, J2S}
J3/M	{J3/M}
J4/M	{J4/M, J4/L, J5M}
J5/M	{J5/M, J4/M, J5/S}
J1/S	{J1/S, J1/L, J1/M}
J2/S	{J2/S, J2/M, J5/S}
J3/S	{J3/S}
J4/S	{J4/S}
J5/S	{J5/S, J2/S, J5/M}

Table 3.6 - Credibility weights

Cells	Classes
J1/L	{J1/L .606, J1/M .357, J1/S .037}
J2/L	{J2/L, 1.00}
J4/L	{J4/L .278, J4/M .722}
J5/L	{J5/L, 1.000}
J1/M	{J1/M .357, J1/L .606, J1/S .037}
J2/M	{J2/M .803, J2S .197}
J3/M	{J3/M, 1.000}
J4/M	{J4/M .381, J4/L .147, J5M .472}
J5/M	{J5/M .517, J4/M .417, J5/S .065}
J1/S	{J1/S .037, J1/L .606, J1/M .357}
J2/S	{J2/S .168, J2/M .683, J5/S .150}
J3/S	{J3/S, 1.000}
J4/S	{J4/S, 1.000}
J5/S	{J5/S .112, J2/S .112, J5/M .789}

Table 3.5 – Revised MLE Estimates (Accidents / Million Departures)

	Jurisdiction 1	Jurisdiction 2	Jurisdiction 3	Jurisdiction 4	Jurisdiction 5
Large	0.554	0.356	N/A	2.889	2.019
Medium	0.554	1.887	4.305	2.804	2.837
Small	0.554	2.216	17.852	9.237	2.850

Unlike schemes that rely on arithmetic functions, this scheme does not force certain relationships to hold across jurisdictions or across size categories. For instance, while the frequency of small airlines is nearly four times that of medium size airlines in Jurisdiction 3, the difference is not nearly as pronounced in other jurisdictions. In fact, in jurisdiction 1, small, medium, and large airlines all have the same accident frequency.

The current classification scheme is one of many that could have been devised using our classification procedure. It should be compared to others to decide which is the most efficient. The notion of efficiency is addressed in [6, p 98].

3.3) Trend

Has the rate of accident changed overtime? If so, how has it changed? The answer to these questions has implications on how the accident rate is projected into the future. Other questions come up as well. Should we look at the change in the rate of accident over the entire cell universe or should we only be concerned with changes within individual cells or within individual classes? Can we even examine the issue of trend independently of that of classification? Should we look at time as one more variable in the classification scheme or do we examine changes over time after the scheme itself has been established? Isn't it possible for some cells to show improvement in frequency overtime while others show deterioration or no change in their accident frequency? Isn't it also possible for the accident rates of two cells or two classes to converge or diverge over time? We do not pretend to have the answers to all these questions. For now, we will look at the question of trend as a one-dimensional problem. In order to measure a trend pattern over time, we first need to specify a model as to how the frequency rate is changing over time. Just as importantly, we have to be able to estimate the parameters of the model and specify the distribution of these parameters.

We now use data from the National Transportation Safety Board shown in table 3.7 below to illustrate how a trend can be estimated and how we can make inferences about the significance of such trend. We will look at a linear and an exponential decay models described, respectively, by equations 1 and 2 below:

$$\hat{\lambda}_i = f_1(y_i) = \alpha + \beta y_i,\ i = 1,2,\cdots,n \qquad \text{Equation 1}$$

$$\hat{\lambda}_i = f_2(y_i) = \alpha' + \beta' e^{\delta'(y_i - y_1)},\ i = 1,2,\cdots,n \text{ where } \delta' \le 0, \alpha' \ge 0 \qquad \text{Equation 2}$$

Other than familiarity perhaps, there is not much rationale for choosing a linear model[12]. We show it here only because its use is so pervasive in casualty actuarial practice. In fact, the NTSB data below will show the potential fallacy of using a linear trend model. The exponential decay model is used commonly in biology and the rationale for its use in biology is applicable in the context of accident frequencies. We can think of the frequency as being made up of two components. The first one, represented by α', is the fixed, intrinsic, or ultimate portion of the accident rate due to a type of error that cannot be eliminated over time whereas the second part $\beta' e^{\delta y_i}$ is the variable portion of the accident rate due to the type of error that changes (decays) over time perhaps as a result of technological advances.

We know that the $\hat{\lambda}_i$'s do not have a constant variance[13]. Therefore, it would be inappropriate to estimate the parameters of equations 1 and 2 by using the ordinary least square function. Since the variance of each $\hat{\lambda}_i$ is inversely proportional[13] to the number of exposures d_i, we instead minimize the following weighted sum of square function: $WSS = \sum_{i=1}^{n} d_i(\hat{\lambda}_i - f_j(y_i))^2$, and we obtain the following weighted least square estimates for models 1 and 2:

[12] Even if a linear model is used, the parameters and their associated errors should not be estimated using simple linear regression. The assumptions of normality and of uniform variance that underlie the simple linear regression model do not hold for the $\hat{\lambda}_i$'s.

[13] $Var(\hat{\lambda}_i) = Var(\frac{x_i}{d_i}) = \frac{1}{d_i^2} Var(x_i) = \frac{\lambda_i d_i}{d_i^2} = \frac{\lambda_i}{d_i}$

Model 1: $\hat{\alpha} = 42.66$ and $\hat{\beta} = -.0212$.

Model 2: $\hat{\alpha}' = 0$, $\hat{\beta}' = .749$, $\hat{\delta}' = -.035$

Figure 3.5 shows the graph for the actual and least square estimates of models 1 and 2. Both models give similar results in the 1982 through 2000 period. However, the models diverge significantly beyond that period. Extrapolation from either model has to be done carefully and should not go out more than a couple of years. However, some situations may call on the modeler to extrapolate over a longer time horizon. For these situations, the exponential decay model may be more appropriate than the linear model. The indications from the linear model are counterintuitive in the long run as they indicate a negative frequency by the year 2018. For the NTSB data, the estimate of the ultimate frequency α' is zero. The indications from the exponential decay model taper off much more slowly and never quite reach zero. Taken at face value, this would be encouraging news for the probability of major accidents in the future.

In appendix B, we show the closed form formula for the weighted least square estimates $\hat{\alpha}$ and $\hat{\beta}$ for model 1 and we also show they are unbiased estimates of α and β. The weighted least square estimates of model 2 are obtained through numerical methods and no closed form formulas are available. What is the distribution of these parameter estimates? Are there statistics that can help us make inferences about the significance of these estimates? What is the error associated with the forecast based on these estimates? We will have to research further for answers to these questions.

Table 3.7 - Major Accident by NTSB Classification for US Air Carriers Operating under CFR 121[14]

Year	Million of Departures	Major Accidents	Frequency per Million Departures
1982	5.35	3	0.56
1983	5.44	4	0.73
1984	5.90	2	0.34
1985	6.31	8	1.27
1986	7.20	4	0.56
1987	7.60	5	0.66
1988	7.72	4	0.52
1989	7.65	8	1.05
1990	8.09	4	0.49
1991	7.81	5	0.64
1992	7.88	3	0.38
1993	8.07	1	0.12
1994	8.24	4	0.49
1995	8.46	3	0.35
1996	8.23	6	0.73

[14]Source: Departures www.ntsb.gov/aviation/Table5.htm; Major Accidents www.ntsb.gov/aviation/Table2.htm.

The NTSB defines a major accident as one that meets any of the following three conditions:
a Part 121 aircraft was destroyed, or
there were multiple fatalities, or
there was one fatality and a Part 121 aircraft was substantially damaged.

The NTSB provides data through the 2001 year. However, starting in 1997, aircrafts with 10 or more seats used in scheduled passenger service began operating under 14 CFR 21. We did not want to analyze the data beyond 1996 as we were not sure whether the inclusion of this new category of aircrafts would distort the indicated trend.

Figure 3.4 – Accident per million departures 1982 – 1996

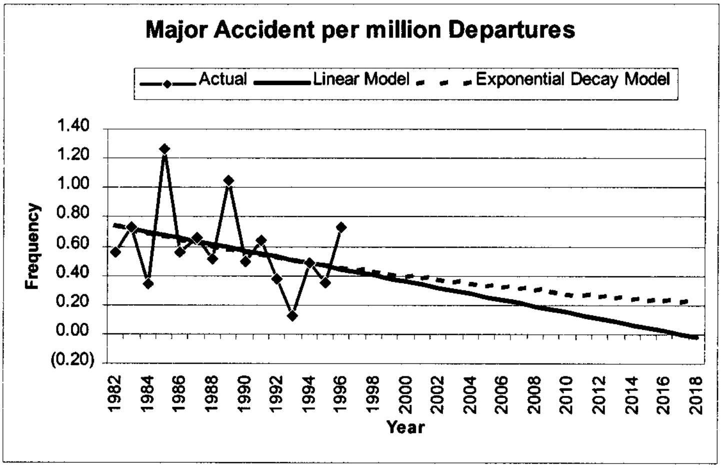

3.4) Modeling the number of aircrafts involved in accidents

Airline catastrophes may be the result of collisions involving several aircrafts. Most accidents involve the failure of a single aircraft. Accidents involving collision of multiple airplanes are relatively rare. However, these types of accidents have occurred and need to be reflected in the simulation model. It would be a mistake not to provide in the model for the possibility of two, three, and perhaps more airplanes being involved in a single collision. We are not talking about collisions triggered by willful acts of sabotage, war, or terrorism. Accidents caused by willful acts will be discussed in section 6. Multiple aircraft collisions can put serious financial strains on the insurers and reinsurers who are responsible for indemnifying the airlines. In addition to the probability distribution of the number of accidents, the modeler needs to specify a conditional distribution for the number of aircrafts involved once there has been an accident. The parameters of this distribution might need to be derived from a fair amount of judgment. A conditional probability distribution table for the number of aircraft involved in an accident is shown in table 3.8 below:

Table 3.8 - Conditional probability for the number of aircrafts involved in an accident

Number of Aircrafts	Probability
1	.97
2	.02
3	.01
4 or more	.00

If an accident is the result of a collision of multiple aircrafts, one also needs to determine the airlines and aircrafts involved. The modeler essentially needs to build yet another conditional probability table laying out the probability of collision between different airlines and aircrafts. Airlines that use a lot of the same airports are more likely to be in a collision than those that use few common airports. Hence, intuitively, these probabilities should be proportional to the proximity of the operation of the airlines and to their relative exposures (say, number of departures). This may seem like a daunting task given the low probability of such events and especially given that the exact identity of the other airlines involved may not be of interest in many applications. However, such table can be greatly simplified by making some broad assumptions. For instance, looking at the probability of two-aircraft collisions for a given airline, one may simply endeavor to compute the conditional probability of collisions involving only aircrafts from that airline. The complement of that probability would be the probability of collisions involving an aircraft from the given airline with one from any other airline.

4) Modeling the cost of catastrophes

The financial costs of airline accidents can be staggering and wide-ranging, affecting individuals, small business entities, corporations, and indeed entire financial markets. The insured portion of these costs is in principle bounded by the parameters of the insurance contract. It is this portion only that we hope to forecast. Here we discuss how to estimate the costs associated with coverage for hull, passenger liability, third party liability, and products liability. However, this model could be used to forecast the costs of other types of coverage such as accident and health, workers compensation, and cargo, for instance, and virtually any financial product where payment is dependent on catastrophic airline accidents.

4.1) Hull Costs

Once an accident has occurred, the hull costs are determined relatively quickly. The insured value of an airline's fleet is pre-determined by the insurance contract. The latter provides for a schedule of insured values for each aircraft in a fleet. The first four columns of table 4.1 below show a typical schedule of aircraft and insured values. This information often does not trickle down to reinsurers, retrocessionaires, or even to some of the smaller primary markets perhaps because it is not used in the rating process. Total fleet value, which is the aggregate of the insured values of individual aircrafts, is usually available but this information is mostly useless in the context of this type of simulation model. In order to accurately forecast the hull cost, the modeler needs to have some idea of the fleet and utilization profile[15] of the airline involved as well as the pre-agreed insured values. It behooves the modeler to make sure this information is obtained. The fleet and utilization profile of an airline is typically public information that can be obtained from the airline's website or from airline industry publications or regulatory agencies. However, insured values need to be obtained through insurance channels. As a substitute for actual insured values, one could estimate the hull cost using the price of a new similar aircraft (ballpark numbers are available from the manufacturers) and factor in a discount based on the aircraft age and configuration. This approach adds a layer of uncertainty in an area where there should be none.

Once the fleet distribution and utilization profile for a given airline is known, the conditional probability of a particular aircraft being involved in an accident can be determined. For instance, the conditional probabilities can simply be calculated as the ratio of each aircraft's projected number of departures to the total number of departures. One may want to factor in the age and model of aircraft in the determination of the conditional probabilities if one believes that these impact the probability of accidents. However, we will work from the basic premise that, for a given airline, the conditional probability of accident for a given aircraft is proportional to its utilization. The fleet and distribution profile, the hull values for a hypothetical airline are shown in table 4.1 below. The conditional probabilities are calculated as indicated above.

[15] The utilization profile refers to the number of hours flown or the number of departures within a period of time for each aircraft within a fleet.

4.1- Fleet, Utilization Profile, and Seating Capacity for a Hypothetical Airline

Aircraft Make and Model	**Registration Number**	**Number of Seats excluding Crew Members**	**Insured Value**	**Projected Utilization (# Departures)**	**Conditional Probability**
A-300-600	XXXXXX	298	118	739	.74%
B-717	XXXXXX	106	40	1,219	1.22%
B-727	XXXXXX	149	60	2,147	2.15%
Total				100,000	100%

4.2) Passenger Liability

There are two important variables in determining the total cost of passenger liability in the event of an accident. The first is the number of passengers involved in an accident while the second is the award per passenger.

4.2.1) Forecasting the number of passengers, survivors, and fatalities involved in an accident

The number of passengers involved in an accident depends on the seating capacity of the aircraft model involved, and the percentage of capacity filled. The seating capacity of each aircraft in a fleet is available in the schedule of aircraft and insured values. If the seating capacity of a given aircraft is not available, one can use the average seating capacity for that specific aircraft model, which can be obtained from many different sources. The actual seating capacity for a given aircraft model may vary based on the specific configuration for that aircraft. The larger the business class and first class sections, the smaller the overall seating capacity. The other important factor in determining the number of passengers involved in an accident is the passenger load factor, which is available through various airline industry publications. The modeler – having determined the aircraft model and therefore the passenger capacity involved in an accident – may use either a fixed or a random passenger load factor to determine the number of passengers on board the aircraft. The modeler may create a simple distribution based on the published load factor and an upper bound of 100%. For instance, if the published load factor is 85%, one might use a uniform distribution with lower and upper bound of 70% and 100%, respectively.

One may also be interested in projecting the number of survivors/fatalities in a given accident. For that, one could look at the survivability statistics for the type of accidents one is investigating and derive a probability distribution for the percentage of survivors/fatalities. That distribution will have a domain bounded by 0% and 100% and will likely be heavily weighted towards these two points. Survivability data can be obtained by reviewing fatality/survival ratios of individual accidents. In table 7.3 below, we show survivability statistics based on data published by the National Transportation Safety Board.

4.2.2) Cost per passenger

The determination of liability cost is more complicated and more involved than that of the number of passengers or injuries involved in an accident. If one tries to simulate the liability cost for each passenger, one has to know the jurisdiction in which compensation will be sought and information regarding the passenger including place of residence, age, marital status, current and projected net worth. Indeed, in the United States for instance, liability awards stemming from a given accident may vary significantly from one passenger to the next. This level of passenger profile detail is not only difficult to obtain but might be unnecessary in most applications. This information will only be relevant if the coverage depends on individual passenger payout such as a layer offering per passenger excess of loss protection. If we focus rather on forecasting the average liability cost per passenger rather than the actual award per passenger, the overriding consideration is the jurisdiction and the applicable laws under which compensation is sought. Such laws are complex, numerous, and constantly evolving. Accidents involving international flights are especially challenging, as even the jurisdiction in which compensation is sought is hard to determine. For instance, under what jurisdiction will compensation be sought in an accident occurring over Canadian land on a flight from New Delhi to New York with a stop in London? Passengers may have different recourses depending on their nationality, the place they purchased their ticket, their final destination on the trip. The modeler has to make some simplifying assumption as to which jurisdiction will be involved in the event of an accident on an international flight. For instance, the modeler may decide that the country on the itinerary with the higher award potential is the country in which suit will be brought. This is especially important for airlines where the liability awards are much smaller in their country of domicile than in some of their other destinations. For these airlines, the liability potential on domestic flights will be significantly lower than the potential on international flights. For accidents involving US airlines wherever occurring, one may simply assume that all suits will be brought in an American court. The modeler needs to have an idea of a carrier's percentage of domestic versus international flights. For international flights, the modeler needs to know the destination profile by country or group of countries. In turn, this information will be used to simulate the itinerary of a flight involved in an accident.

Does the modeler then need to know the distribution of the average liability award in every possible jurisdiction or country? This would be a daunting task even for the most industrious modeler. The modeler may instead look to group countries where the tort and compensation systems are similar. For instance, a group might be comprised of countries in the European Community, another of Mercosur countries[16]. Also, instead of looking at the distribution of actual average award, the modeler may choose instead to look at the average award as a ratio to the median income or the income per capita in a country. Assuming that the modeler can come up with such groups, there remains the challenge of using the historical data to come up with the average award distribution. Since liability claims can take an inordinate amount of time to settle, an average award distribution will need to be built largely on case estimates, the accuracy of which won't be known for a long time. For those claims that have already been settled, inflation, changes in law, voluntary agreements, and statutes may render them less relevant for the purpose of projecting the cost of future claims. The upshot of all this is that an average liability award distribution will involve significant judgment on the part of both the modeler and others. Once groups of countries have been defined, the modeler may try to find a distribution that fits the actual data adjusted for inflation and for past and expected future changes. Similarly to the frequency portion of the model, should the modeler devise some statistical tests to help him decide whether the distribution of liability awards (as a percentage of, say, median income) for the various groups are indeed dissimilar? For instance, upon close analysis, it may turn out that the average award potential in the European Community is not dissimilar to that in the Mercosur countries.

4.3) Third Party Liability

Third party liability costs are even more difficult to forecast. The range of scenarios for third party liability is obviously much wider than that for passenger liability. Also, relatively few commercial airline accidents result in injury or property damage to third parties. This is perhaps because the location of airports and the air routes have tended to steer airplanes away from populated areas. However, these events have occurred and need to be considered in the simulation model. One approach might be to look at the passenger and third party liability together. So instead of looking at the distribution of the average passenger liability per passenger as suggested in the preceding section, one would look at the average total liability per passenger. The tail of that distribution would be a lot more skewed than that of the average passenger liability award. Another approach is to estimate the number or percentage of accidents that will result in third party damage and to estimate the cost of such liabilities separately. This approach is better at allowing the modeler to factor in extreme scenarios. For instance, the modeler might include a scenario where, as a result of a midair collision, two jumbo airplanes plow onto a crowded area destroying life and property. There, considerable judgment might be used to determine the likelihood of different scenarios.

4.4) Products Liability

Defendants in lawsuits stemming from aircraft accidents include not only the airlines but also aircraft and parts manufacturers as well as other parties involved in the operation of the airline. Also, the airlines themselves can try to recoup losses by suing other parties not necessarily named in a suit. For this reason, aircraft and parts manufacturers require products liability to shield them from such suits. To understand the products liability exposure in the context of this simulation model, information about manufacturers and suppliers of engine, navigation equipment, electrical system and other components has to be collected for each insured aircraft. The identity of the aircraft manufacturer itself should be obvious. In the event of an accident, we would then have a list of potential defendants. We have

[16] Mercosur countries, as of the time of this writing, are made up of Agrentina, Brazil, Uruguay, and Paraguay.

already shown how the total liability from a given accident can be estimated. We then need to allocate that liability between airline operators and the product manufacturers. The actual allocation of liability will depend very much on what is determined to be the cause of accident, and, furthermore, may vary from one jurisdiction to the next. This is a very difficult area in need of much research. Considerable judgment or simplification might be used to come up with an allocation. In order to figure out the product liability exposure of a given manufacturer, we would then accumulate its exposure over the entire universe of airline operators.

5) Validation

Before using any model to forecast, the modeler needs to make sure that the model's assumptions are reasonable. The particular model we have presented makes many assumptions, one nested inside another. To develop any sense of how well the model will forecast the future, one can look at how well the model would have predicted past experience. Let's say that data from 1980 through 2000 is available. The modeler may endeavor to see how well the simulation model would have projected 1991 based on data through 1990, 1992 based on data through 1991, and so on, thus obtaining a comparison with actual data for ten years. The validation should be done in stages, starting with a look at the number of accidents, the number of passengers involved in accidents, the number of fatalities and injuries, and the insurance costs in that order. Doing the validation in stages allows one to identify where in the simulation process a bias may be occurring and to make the necessary adjustments. The results of the simulation will be a probability distribution for the projected variables similar to the one shown in table 7.5 of section 7 below. Focusing on the number of fatalities for instance, one would compare the actual number of fatalities in 1991, 1992 and subsequent years with the distribution predicted for each of these years by the simulation model. If, on one hand, the actual number of fatalities for the ten-year sample (1991 through 2000) looks like a random draw from each of the predicted distribution, this would tend to validate the simulation model. If, on the other hand, the actual number of fatalities for the ten-year sample tends to fall systematically either to the right of, say, the 95th percentile or to the left of, say, the 5th percentile of the predicted distributions, this would be a strong indication that the model's projections are biased.

6) Terrorism

The audacity and severity of the events of September 11th caught most insurance professionals off guard. Reportedly, aviation underwriters had been "giving away" for free the coverage for terrorist acts. In light of the devastating potential of terrorist acts, actuaries and underwriters have since been scrambling to put a price on terrorism coverage. While some pundits will offer nothing more than their eloquent prose enlightening us with revelations such as "...the risk [of terrorism] is real", underwriters and actuaries are left with the unenviable task of putting a price on the risk of terrorism. Perhaps, in no other area will the actuary need to use all available sources of information and rely on the expertise of others in order to try to quantify the risk of terrorism. Many of the assumptions we have made in relation to accidental airline catastrophes certainly don't apply to crashes occurring as a result of willful acts. We know that terrorist acts are neither random nor are they uncorrelated. This contradicts the assumptions implicit in our use of the Poisson distribution. We also know that history is perhaps a poor guide for figuring out future acts. The risk of terrorism is highly fluid as our geopolitical landscape changes constantly, and as airlines and law enforcement authorities learn how to better protect the public from such acts. In our preceding discussion, we rely extensively on historical data to derive expected frequencies. We could not do the same with terrorism although a look at the history can be instructive. Table 6.2 shows the number of hijackings perpetrated against US and foreign airlines, respectively, from 1970 through 2000. The sharp drop in the number of hijackings against US airlines, with none recorded between 1992 and 2000, testifies perhaps to the success US airlines and authorities have had in deterring and preventing such acts. The events of September 11th, 2001 serve as a staunch reminder that the probability of such acts is never quite zero. Although, we have focused here on hijackings, they are by no means the only terrorist threat facing airlines. We should also keep in mind that not all hijackings have resulted in death, injury, or destruction of property as people have sought to hijack airplanes for a variety of reasons. Hijackings need not be the result of some political conflicts. For instance, mentally deranged individuals with no apparent political motives have hijacked airplanes.

Although individual airlines may have their own security procedures, the modeler can work from the assumption that the risk of terrorist acts against an airline depends primarily on the level of security in the airports and countries where the airline operates. Actuaries could work with security experts in developing a grouping system to rate airports' and countries' potential for terrorist acts. The level of conflict in a country and its surrounding regions as well as a country's will and ability to effectively fight terrorism need to be factored in such a grouping. Actuaries could then try to formulate a probability of an airline being hit by a terrorist for each category in a grouping. For instance, table 6.1 below shows a hypothetical two-dimensional box which groups airports based on their level of security and the existence of a terrorist threat around them. The probability of an airline's being hit by a terrorist act would be calculated as the weighted average of the probabilities of each airport where it has exposure.

Table 6.1 – Probability (odds in 1 million) of a Terrorist Act in a 12 Month Period

Security / Terror Threat	Impenetrable	Strict	Adequate	Lax	Non-existent
Constant	10	100	250	1,000	10,000
Potential	4	50	200	800	9,500
Some	2	20	150	700	8,000
None	1	5	10	500	6,000

As we mentioned before, these probabilities are dynamic and should change as conflicts evolve or new ones emerge, as new information comes to light, and as new acts of terrorism are committed or attempted. This implies that the price for such coverage will, at least in theory, be dynamic, changing as the risk of terrorist acts is continuously reassessed.

Table 6.2 - Hijackings

Year	U.S.	Foreign	Year	U.S.	Foreign	Year	U.S.	Foreign
1970	25	49	1981	7	24	1992	0	12
1971	25	30	1982	9	23	1993	0	31
1972	26	30	1983	17	15	1994	0	23
1973	2	20	1984	5	21	1995	0	9
1974	3	17	1985	4	22	1996	0	14
1975	6	13	1986	2	5	1997	0	10
1976	2	14	1987	3	5	1998	0	9
1977	5	26	1988	1	10	1999	0	11
1978	7	16	1989	1	14	2000	0	20
1979	11	13	1990	1	39			
1980	21	18	1991	1	23			

Sources: 1970-1998 U.S. Department of Transportation, Federal Aviation Administration, Criminal Acts Against Civil Aviation - 1998, Charts and Graph; 1999-2000 http://cas.faa.gov/crimacts/doc/crim2000.doc Appendices A&B

7) A simplified application of a simulation model

A simulation model has many applications for assessing the costs of insurance coverages and other financial instruments affected by airline catastrophes. Here, we present an example where we look at the costs of a cover that pays for the full insured value of a destroyed or damaged aircraft as well as $50,000 per fatality and $100,000 per injured passenger for a hypothetical group of airlines for the 2003 year. This coverage excludes acts of war and terrorism. The information and assumptions are set in the tables 7.1 through 7.4. The results of the simulation are presented in table 7.5.

Table 7.1 - Fleet, Utilization Profile, and Seating Capacity

Aircraft Type	Count	Seats	Insured Value (MM)	# of Departures	Prob
Airbus Industrie A300-600	79	298	118	58,390	0.69%
Airbus Industrie A300B2/B4	19	298	118	6,165	0.07%
Airbus Industrie A310	44	249	92	22,253	0.26%
Airbus Industrie A319	137	125	52	148,736	1.77%
Airbus Industrie A320	227	172	55	287,301	3.42%
Airbus Industrie A380	6	600	250	4,729	0.06%
Avro RJ Avroliner	36	70	26	71,171	0.85%
BAE SYSTEMS (HS) 146	18	94	40	47,420	0.56%
Boeing (McDonnell-Douglas) DC-10	239	264	110	123,748	1.47%
Boeing (McDonnell-Douglas) DC-8	194	146	60	92,469	1.10%
Boeing (McDonnell-Douglas) DC-9	430	115	50	684,705	8.15%
Boeing (McDonnell-Douglas) MD-11	66	325	150	41,923	0.50%
Boeing (McDonnell-Douglas) MD-80	670	155	60	1,056,052	12.57%
Boeing (McDonnell-Douglas) MD-90	21	163	60	34,382	0.41%
Boeing 717	31	106	40	37,779	0.45%
Boeing 727	729	167	60	715,326	8.51%
Boeing 737 (CFMI)	779	149	60	1,672,505	19.90%
Boeing 737 (JT8D)	248	149	60	518,652	6.17%
Boeing 737 (NG)	303	149	60	372,020	4.43%
Boeing 747 Classic	138	472	215	66,097	0.79%
Boeing 747-400	73	544	211	37,837	0.45%
Boeing 757	589	214	85	728,728	8.67%
Boeing 767	328	251	110	274,901	3.27%
Boeing 777	98	373	190	60,625	0.72%
Bombardier (Canadair) CRJ Regional Jet	273	50	25	543,579	6.47%
Embraer ERJ-135	53	36	14	64,287	0.76%
Embraer ERJ-145	169	50	20	274,355	3.26%
Fairchild/Dornier 328JET	22	30	13	20,837	0.25%
Fokker 100	123	113	50	234,759	2.79%
Fokker F.28	22	85	40	53,688	0.64%
Lockheed L-1011 TriStar	81	280	110	48,412	0.58%
Total	6,245	1,108,002	450,332	8,403,831	

Table 7.2 - Projected Exposures, Frequencies, and Passenger Loads

Year	2003
Projected Departures	8,918,213
Projected Average Passenger Load	0.65
Expected Frequency per million departures	0.45

- Distribution of survival ratio based on Beta distribution fitted to data in table 7.3 below:

Table 7.3 - Empirical Survival percentages from a sample of 27 accidents

Passengers	Survivors	Fatalities	% of Survivors
31	0	31	0.0%
25	0	25	0.0%
132	0	132	0.0%
68	0	68	0.0%
110	0	110	0.0%
230	0	230	0.0%
155	1	154	0.6%
71	1	70	1.4%
163	29	134	17.8%
57	20	37	35.1%
51	24	27	47.1%
296	185	111	62.5%
82	54	28	65.9%
89	67	22	75.3%
44	36	8	81.8%
108	94	14	87.0%
145	134	11	92.4%
33	32	1	97.0%
149	147	2	98.7%
142	142	0	100.0%
39	39	0	100.0%
23	23	0	100.0%
40	40	0	100.0%
102	102	0	100.0%
292	292	0	100.0%
62	62	0	100.0%
2,739	***1,524***	***1,215***	***55.6%***

Source: National Transportation Safety Board, "Survivability of Accidents Involving Part 121 U.S. Air Carrier Operations, 1983 Through 2000" Safety Report NTSB/SR-01/01, table 4, p. 14. http://www.ntsb.gov/Publictn/A_Stu.htm.

Table 7.4 - Conditional Probability for the Number of Aircrafts Involved in an Accident

Number of Aircrafts	Probability
1	.970
2	.029
3	.001
4 or more	.000

Additional Assumptions

- Destroyed aircraft is immediately replaced with similar aircraft.
- Projected departures, passenger loads and frequencies are fixed. More realistically, these should be allowed to vary.
- Collision occurs between aircrafts within group.
- Selections of parameters only loosely based on real data

Table 7.5 - Simulation Results from 5000 iterations

Variables	Accident Count	Aircraft Count	Passger Count	Injured Count	Death Count	Hull Cost (MM)	Passger Cost (MM)	Total Cost (MM)
Best Case	0	0	0	0	0	0	0	0
Wrst Case	13	14	1,790	1,015	1,012	1,131	1,376	2,507
Expected	3.59	3.69	375	197	179	233	286	519
Std Dev	1.91	1.99	230	154	144	139	182	315
5%	1	1	58	0	0	50	42	95
10%	1	1	109	18	8	60	76	149
15%	2	2	144	43	30	95	105	204
20%	2	2	176	65	51	115	128	248
25%	2	2	204	81	70	120	150	283
30%	2	3	231	97	86	145	172	322
35%	3	3	259	114	103	170	194	363
40%	3	3	287	131	121	180	214	398
45%	3	3	317	148	135	195	236	434
50%	3	4	347	167	150	216	258	474
55%	4	4	377	184	168	230	282	518
60%	4	4	407	209	190	252	305	559
65%	4	4	440	230	212	270	333	604
70%	4	5	474	254	233	290	364	652
75%	5	5	509	281	258	315	396	714
80%	5	5	554	314	287	345	430	772
85%	6	6	614	355	325	375	473	845
90%	6	6	681	407	371	415	531	932
95%	7	7	792	488	447	485	627	1,089

The distribution presented in table 7.5 above shows the variability in the loss process but does not incorporate parameter error. We have assumed for instance that the claim process follows a Poisson distribution with an expected frequency per million departures of .45. In reality, we will never know the true expected frequency of such distribution or the exact form of the distribution for that matter. At best, we will have an estimate of the frequency with an error margin. The conditional probability distribution for the number of aircrafts involved in an accident, the distribution of the passenger load factors and of the survival ratios are all subject to parameter error. One way to incorporate the parameter uncertainty into the results would be to allow the parameters to vary according to some distribution. A different but simpler approach might be to look at different combinations of parameter estimates and to run the simulation for each combination.

Applications for this type of simulation model extend much beyond the type of examples we have just presented. We have only scratched the surface of the possibilities that a simulation approach offers. The CAS's literature is replete with articles on how simulation models could be used to structure and price reinsurance products. Please see [1], [2], and [4] for a sample of such articles. Obviously the more detailed information available to the modeler, the more accurate the projections will be and the more specific the applications will be for this type of model. Modelers have to weigh the trouble of gathering the additional data against the additional accuracy and flexibility that would be gained.

8) Final Thoughts

Compared to a traditional experience rating approach, a simulation approach promises much more in terms of accuracy and range of applications. For instance, the pricing of reinsurance contracts that feature loss triggers, aggregate limits and deductibles, contingent profit will be more readily handled through simulation than through experience rating. The body of available actual experience would often not suffice to test the multitude of scenarios that can present themselves under such contracts. For the cover we introduced in the preceding section, assume we were interested in pricing an aggregate layer providing 1.0B in limit in excess of a 1.0B retention. According to the aggregate loss distribution in table 7.5 above, there is about a 5% chance that actual aggregate losses would exceed the 1.0B retention in any one year. Looking at the actual experience, over a five to ten year period, may not reveal any losses in the layer. Even when there would be losses in the layer, they may not have much predictive value.

Many of the benefits of using a stochastic simulation approach in the evaluation of property catastrophe extend to catastrophic airline exposures as well. Rade T. Musulin, in an article titled "Issues in the Regulatory Acceptance of Computer Modeling for Property Insurance Ratemaking", [3, p 354] lists the comprehensibility of prices, rational behavior, fair pricing, reduced information risk, and stable pricing as benefits of the improved estimates provided by a simulation model in evaluating property catastrophe exposures. It should be readily apparent how the use of a simulation approach in evaluating airline catastrophe would enhance the comprehensibility of prices, reduce information risk, and promote stable pricing.

Ultimately, whether simulation models gain acceptance in the commercial aviation realm will depend on whether the perceived benefits outweigh the additional effort required in implementing such models. In a sense, the widespread use and acceptance of property catastrophe modeling should have already paved the way for the use of simulation models not only in commercial aviation but also in other lines such as surety, credit, and workers compensation. These days, insurers and reinsurers have the means necessary to keep large amount of information about their property exposures at a zip code, and sometimes, finer level. More importantly, this level of detailed information is accepted as a normal course of doing business in the property catastrophe lines. The task of gathering information on individual aircrafts is relatively small due to the limited number of commercial airlines servicing the world and also the limited number of aircrafts they operate. For instance, the current fleet of US-domiciled airlines consists of less than 6,500 aircrafts of roughly 30 different models. Furthermore, two companies, Boeing and Airbus, manufacture 85% of these aircrafts. Once one goes through the trouble of tallying that information, subsequent updates should be relatively simple as airlines do not change their fleet drastically overnight.

We have left open a number of issues including inferences about the parameters of the exponential decay trend model. We also think that a substantial amount of work has to be done to make realistic projections for third party liability exposures and in the allocation of liability between airline operators and manufacturers. Finally, we have barely broached the issue of terrorism. However, we are confident that there will be a wealth of papers addressing these issues in a much more comprehensive fashion.

REFERENCES

[1] Bernes, Regina M, "Reinsurance Contracts with a Multi-Year Aggregate Limit", CAS Forum, Spring 1997.

[2] Daino, Robert A; Thayer, Charles A, "Comparing Reinsurance Programs – A Practical Actuary's System", CAS Forum, Spring 1997.

[3] Musulin, Rade, T, "Issues in the Regulatory Acceptance of Computer Modeling for Property Insurance Ratemaking", Journal of Insurance Regulation, Spring 1997.

[4] Papush, Dmitry E, "A Simulation Approach in Excess Reinsurance Pricing", CAS Forum, Spring 1997.

[5] Patzor, Andy, "Bumpy Ride", The Wall Street Journal (September 19, 2000), p A1.

[6] Salam, Romel G, "Reinventing Risk Classification – A Set Theory Approach", CAS Forum, Winter 2002.

Appendix A

Risk Classification Procedure

Following the procedure introduced in "Risk Classification – A Set Theory Approach" and at the presentation of the paper at the winter 2002 meeting, we take the following steps:

i. We calculate the $\hat{R}_0$ values for all pairs of adjacent cells[1] as shown in exhibit 1. In figure A.1 below, we show the $\hat{R}_0$ values between the cell Jurisdiction 1/Large (J1/L) and the cells that are adjacent to it.

Figure A.1 - $\hat{R}_0$ between large airlines in jurisdiction 1 (J1/L) and those in adjacent cells

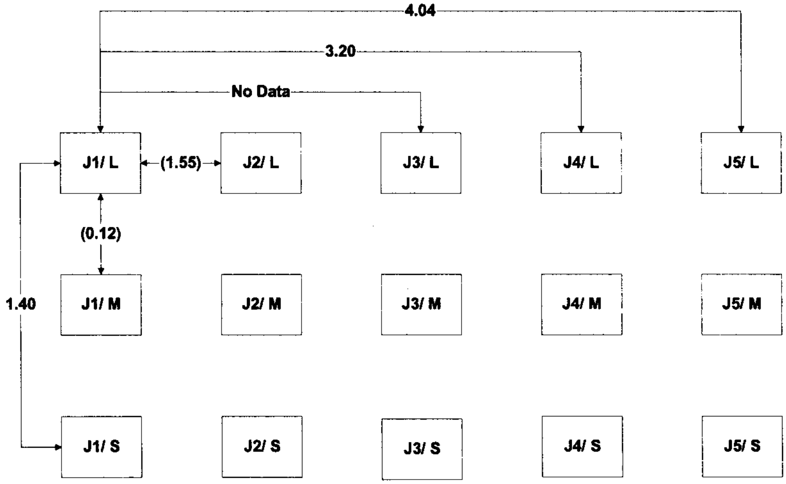

Those adjacent cells for which the $\hat{R}_0$ values fall within the interval (-1.65,1.65) are said to be compatible. All other pairs of cells are said to be incompatible. Exhibit 2 shows the compatibility relationship for all pair of cells. Below, in figure A.2, we answer the question of compatibility for large airlines in jurisdiction 1.

[1] Two cells are said to be adjacent if they have at least one common risk characteristic [6, p.89].

Figure A.2 – Are cells compatible with J1/L (before validation)?

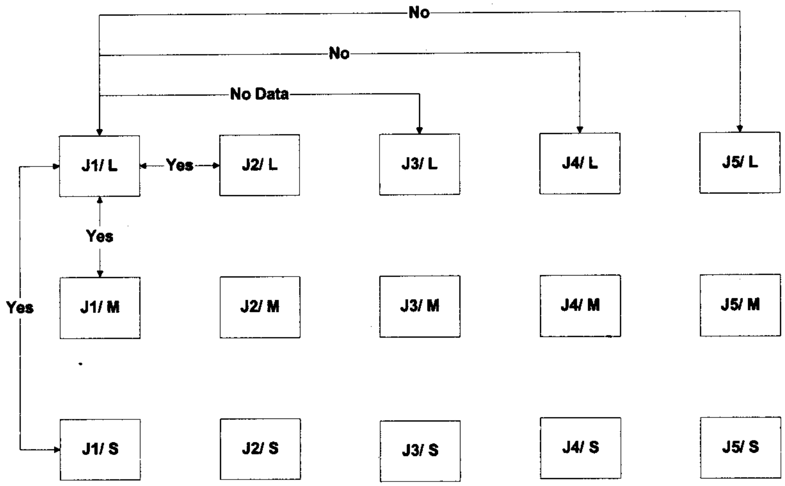

ii. The values of $\hat{R}_0$ shown in exhibit 2 and in figure A.2 may have been the result of a chance occurrence or some oddity in the data and may not reflect the true relationship between cells. Rather than relying on just one drawing of $\hat{R}_0$, we repeat the calculation of the $\hat{R}_0$ values for 1,000 randomly selected samples from each cell in order to validate the compatibility relationships. A sample consists of a random draw of between 50% and 85% of the exposures (or airlines) within a cell. If the number of times $\hat{R}_0$ falls in the interval (-1.65, 1.65) for a given pair of cells is large (greater or equal to 875), then, compatibility is validated for that pair. Exhibits 3 and 4 show the number of times $\hat{R}_0$ falls in the interval for all pairs of adjacent cells and whether these cells are deemed compatible, respectively. This information is shown for the cell Jurisdiction 1/Large in figures A.3 and A.4 below.

Figure A.3 - Number of times out of 1,000 trials $\hat{R}_0$ falls in (-1.65, 1.65) interval

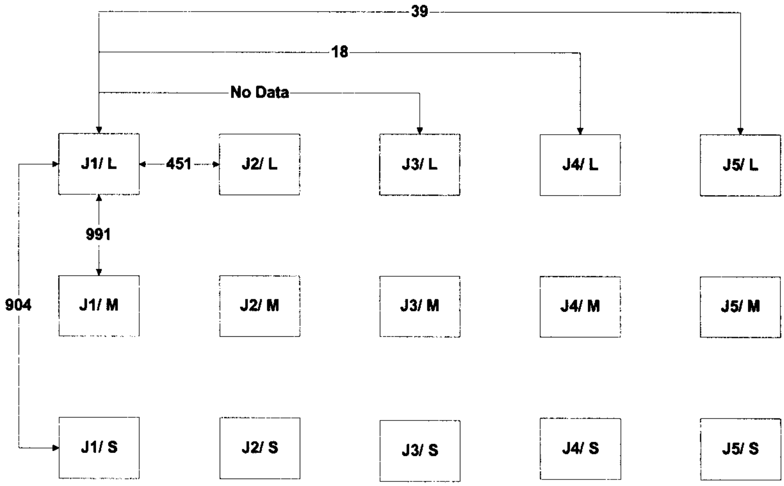

Figure A.4 – Are cells compatible with J1/L (after validation)?

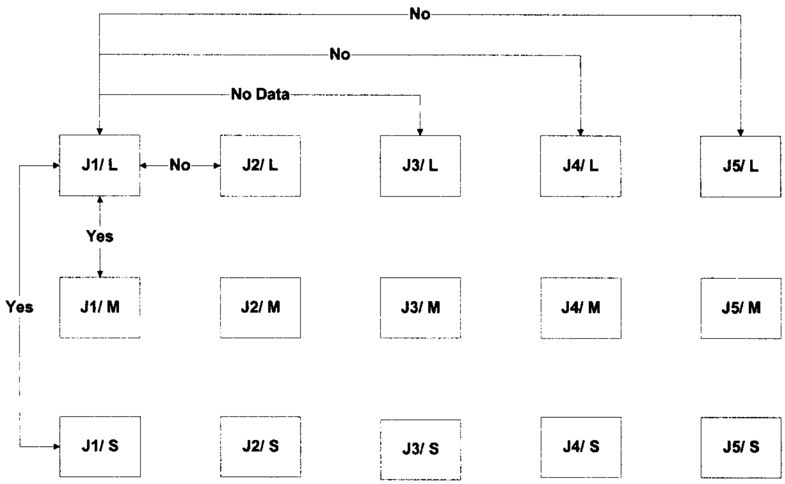

iii. Each cell defines a class made up of these cells that are compatible with it, including the cell itself. In tables A.1 and A.2, respectively, we show the classes defined by each cell and the credibility weights assigned to each cell in a class. The class defined by large airlines in jurisdiction 1 consists of the following three cells: {Jurisdiction 1/Large, Jurisdiction 1/Medium, and Jurisdiction 1/Small}

iv. The revised MLE estimates for each cell is the weighted average of the MLE estimates of the cells in its class where the credibility weights are the exposures in each cell relative to the total exposures for the class.

Table A.1 – Calculation of Revised MLE Estimates for Jurisdiction 1/Large

	Jurisdiction 1 Large	Jurisdiction 1 Medium	Jurisdiction 1 Small	**Total/Weighted Average**
Initial MLE	0.527	.507	1.443	**.554**
Departures	56.93	33.50	3.46	**214.74**
Weights	0.61	0.36	0.04	**1.00**

Table A.2 – Revised MLE Estimates

	Jurisdiction 1	Jurisdiction 2	Jurisdiction 3	Jurisdiction 4	Jurisdiction 5
Large	0.554	0.356	N/A	2.889	2.019
Medium	0.554	1.887	4.305	2.804	2.837
Small	0.554	2.216	17.852	9.237	2.850

Appendix B

Derivation of parameter estimates for a linear trend Model

Let X_i be a Poisson distributed random variable for year y_i, $i = 1,2,\cdots,n$, with mean $\lambda_i d_i$, where d_i represents the number of exposures associated with year y_i.

We posit the following linear relationship between the λ_i's :

$$\lambda_i = \alpha + \beta y_i,\ i = 1,2,\cdots,n$$

Let $\Lambda_i = \dfrac{X_i}{d_i}$ be the random variable representing the maximum likelihood estimate of λ_i.

$E(\Lambda_i) = E(\dfrac{X_i}{d_i}) = \lambda_i = \alpha + \beta y_i$

Also, $Var(\Lambda_i) = \dfrac{\lambda_i}{d_i}$

We denote by x_i and $\hat{\lambda}_i = \dfrac{x_i}{d_i}$ the realizations of the random variables X_i and Λ_i, respectively.

We define the weighted least square error function:

$$WLSE = \sum_{i=1}^{n} d_i(\hat{\lambda}_i - \lambda_i)^2 = \sum_{i=1}^{n} d_i(\hat{\lambda}_i - \alpha - \beta y_i)^2$$

Let $\hat{\alpha}$ and $\hat{\beta}$ represent the values of α and β that minimize the weighted least square function. $\hat{\alpha}$ and $\hat{\beta}$ are obtained as follows:

$$\frac{dWLSE}{d\alpha} = -2\sum_{i=1}^{n} d_i(\hat{\lambda}_i - \alpha - \beta y_i) = 0 \qquad (1)$$

$$\frac{dWLSE}{d\beta} = -2\sum_{i=1}^{n} d_i y_i(\hat{\lambda}_i - \alpha - \beta y_i) = 0 \qquad (2)$$

Solving these two equations simultaneously yields:

$$\hat{\beta} = \frac{\sum_{i=1}^{n} d_i({}_w y - y_i)\hat{\lambda}_i}{\sum_{i=1}^{n} d_i[{}_w y^2 - {}_w(y^2)]} \text{ and } \hat{\alpha} = {}_w\hat{\lambda} - \beta\,{}_w y, \text{ where}$$

$$_w y = \frac{\sum_{i=1}^{n} d_i y_i}{\sum_{i=1}^{n} d_i}, \; _w y^2 = (_w y)^2, \; _W(y^2) = \frac{\sum_{i=1}^{n} d_i y_i^2}{\sum_{i=1}^{n} d_i}, \text{ and } _W\hat{\lambda} = \frac{\sum_{i=1}^{n} d_i \hat{\lambda}_i}{\sum_{i=1}^{n} d_i}$$

Let $\omega_i = \dfrac{d_i(_W y - y_i)}{\sum_{i=1}^{n} d_i[_W y^2 - _W(y^2)]}$, we rewrite $\hat{\beta} = \sum_{i=1}^{n} \omega_i \hat{\lambda}_i$.

$\hat{\beta}$ is the realization of the random variable $B = \sum_{i=1}^{n} \omega_i \Lambda_i$

$$E(B) = \sum_{i=1}^{n} \omega_i E(\Lambda_i) = \sum_{i=1}^{n} \omega_i \lambda_i = \sum_{i=1}^{n} \omega_i E(\alpha + \beta y_i) = E(\alpha)\sum_{i=1}^{n} \omega_i + E(\beta)\sum_{i=1}^{n} \omega_i y_i = E(B) = \beta$$

Therefore $\hat{\beta}$ is an unbiased estimate of β.

Also, $Var(B) = \sum_{i=1}^{n} \omega_i^2 Var(\Lambda_i) = \sum_{i=1}^{n} \omega_i^2 \dfrac{\lambda_i}{d_i}$

Exhibit 1

Calculation of $\hat{R}_0$ between adjacent cells

	J1/L	J2/L	J4/L	J5/L	J1/M	J2/M	J3/M	J4/M	J5/M	J1/S	J2/S	J3/S	J4/S	J5/S
J1/L	-	1.55	(3.20)	(4.04)	0.12	//////	//////	//////	//////	(1.40)	//////	//////	//////	//////
J2/L	(1.55)	-	(3.42)	(4.61)	//////	(3.03)	//////	//////	//////	//////	(2.13)	//////	//////	//////
J4/L	3.20	3.42	-	1.25	//////	//////	//////	0.34	//////	//////	//////	//////	(2.98)	//////
J5/L	4.04	4.61	(1.25)	-	//////	//////	//////	//////	(1.26)	//////	//////	//////	//////	(1.39)
J1/M	(0.12)	//////	//////	//////	-	(2.60)	(4.93)	(4.69)	(5.05)	(1.42)	//////	//////	//////	//////
J2/M	//////	3.03	//////	//////	2.60	-	(3.00)	(1.75)	(1.71)	//////	(0.88)	//////	//////	//////
J3/M	//////	//////	//////	//////	4.93	3.00	-	1.68	1.84	//////	//////	(5.14)	//////	//////
J4/M	//////	//////	(0.34)	//////	4.69	1.75	(1.68)	-	0.14	//////	//////	//////	(3.31)	//////
J5/M	//////	//////	//////	1.26	5.05	1.71	(1.84)	(0.14)	-	//////	//////	//////	//////	(0.91)
J1/S	1.40	//////	//////	//////	1.42	//////	//////	//////	//////	-	(1.00)	(6.30)	(3.91)	(1.67)
J2/S	//////	2.13	//////	//////	//////	0.88	//////	//////	//////	1.00	-	(5.48)	(2.97)	(0.74)
J3/S	//////	//////	//////	//////	//////	//////	5.14	//////	//////	6.30	5.48	-	2.73	4.73
J4/S	//////	//////	2.98	//////	//////	//////	//////	3.31	//////	3.91	2.97	(2.73)	-	2.17
J5/S	//////	//////	//////	1.39	//////	//////	//////	//////	0.91	1.67	0.74	(4.73)	(2.17)	-

Exhibit 2

Are cells compatible (before validation)?

	J1/L	J2/L	J4/L	J5/L	J1/M	J2/M	J3/M	J4/M	J5/M	J1/S	J2/S	J3/S	J4/S	J5/S
J1/L	Yes	Yes	No	No	Yes	//////	//////	//////	//////	Yes	//////	//////	//////	//////
J2/L	Yes	Yes	No	No	//////	No	//////	//////	//////	//////	No	//////	//////	//////
J4/L	No	No	Yes	Yes	//////	//////	//////	Yes	//////	//////	//////	//////	No	//////
J5/L	No	No	Yes	Yes	//////	//////	//////	//////	Yes	//////	//////	//////	//////	Yes
J1/M	Yes	//////	//////	//////	Yes	No	No	No	No	Yes	//////	//////	//////	//////
J2/M	//////	No	//////	//////	No	Yes	No	No	No	//////	Yes	//////	//////	//////
J3/M	//////	//////	//////	//////	No	No	Yes	No	No	//////	//////	No	//////	//////
J4/M	//////	//////	Yes	//////	No	No	No	Yes	Yes	//////	//////	//////	No	//////
J5/M	//////	//////	//////	Yes	No	No	No	Yes	Yes	//////	//////	//////	//////	Yes
J1/S	Yes	//////	//////	//////	Yes	//////	//////	//////	//////	Yes	Yes	No	No	No
J2/S	//////	No	//////	//////	//////	Yes	//////	//////	//////	Yes	Yes	No	No	Yes
J3/S	//////	//////	//////	//////	//////	//////	No	//////	//////	No	No	Yes	No	No
J4/S	//////	//////	No	//////	//////	//////	//////	No	//////	No	No	No	Yes	No
J5/S	//////	//////	//////	Yes	//////	//////	//////	//////	Yes	No	Yes	No	No	Yes

Exhibit 3

Number of times $\hat{R}_0$ falls in (-1.65,1.65) interval

	J1/L	J2/L	J4/L	J5/L	J1/M	J2/M	J3/M	J4/M	J5/M	J1/S	J2/S	J3/S	J4/S	J5/S
J1/L	-	451	18	39	991	//////	//////	//////	//////	904	//////	//////	//////	//////
J2/L	451	-	0	18	//////	38	//////	//////	//////	//////	393	//////	//////	//////
J4/L	18	0	-	824	//////	//////	//////	966	//////	//////	//////	//////	129	//////
J5/L	39	18	824	-	//////	//////	//////	//////	791	//////	//////	//////	//////	861
J1/M	991	//////	//////	//////	-	189	1	4	10	905	//////	//////	//////	//////
J2/M	//////	38	//////	//////	189	-	170	642	760	//////	999	//////	//////	//////
J3/M	//////	//////	//////	//////	1	170	-	583	449	//////	//////	10	//////	//////
J4/M	//////	//////	966	//////	4	642	583	-	904	//////	//////	//////	75	//////
J5/M	//////	//////	//////	791	10	760	449	904	-	//////	//////	//////	//////	932
J1/S	904	//////	//////	//////	905	//////	//////	//////	//////	-	809	0	0	413
J2/S	//////	393	//////	//////	//////	999	//////	//////	//////	809	-	1	127	976
J3/S	//////	//////	//////	//////	//////	//////	10	//////	//////	0	1	-	315	20
J4/S	//////	//////	129	//////	//////	//////	//////	75	//////	0	127	315	-	455
J5/S	//////	//////	//////	861	//////	//////	//////	//////	932	413	976	20	455	-

Exhibit 4

Are cells compatible (after validation)?
Cut off point 875

	J1/L	J2/L	J4/L	J5/L	J1/M	J2/M	J3/M	J4/M	J5/M	J1/S	J2/S	J3/S	J4/S	J5/S
J1/L	Yes	No	No	No	Yes	//////	//////	//////	//////	Yes	//////	//////	//////	//////
J2/L	No	Yes	No	No	//////	No	//////	//////	//////	//////	No	//////	//////	//////
J4/L	No	No	Yes	No	//////	//////	//////	Yes	//////	//////	//////	//////	No	//////
J5/L	No	No	No	Yes	//////	//////	//////	//////	No	//////	//////	//////	//////	No
J1/M	Yes	//////	//////	//////	Yes	No	No	No	No	Yes	//////	//////	//////	//////
J2/M	//////	No	//////	//////	No	Yes	No	No	No	//////	Yes	//////	//////	//////
J3/M	//////	//////	//////	//////	No	No	Yes	No	No	//////	//////	No	//////	//////
J4/M	//////	//////	Yes	//////	No	No	No	Yes	Yes	//////	//////	//////	No	//////
J5/M	//////	//////	//////	No	No	No	No	Yes	Yes	//////	//////	//////	//////	Yes
J1/S	Yes	//////	//////	//////	Yes	//////	//////	//////	//////	Yes	No	No	No	No
J2/S	//////	No	//////	//////	//////	Yes	//////	//////	//////	No	Yes	No	No	Yes
J3/S	//////	//////	//////	//////	//////	//////	No	//////	//////	No	No	Yes	No	No
J4/S	//////	//////	No	//////	//////	//////	//////	No	//////	No	No	No	Yes	No
J5/S	//////	//////	//////	No	//////	//////	//////	//////	Yes	No	Yes	No	No	Yes

Exhibit 5

Classes defined by each cell

Cells	Classes
J1/L	{J1/L, J1/M, J1/S}
J2/L	{J2/L}
J4/L	{J4/L, J4/M}
J5/L	{J5/L}
J1/M	{J1/M, J1/L, J1/S}
J2/M	{J2/M, J2S}
J3/M	{J3/M}
J4/M	{J4/M, J4/L, J5M}
J5/M	{J5/M, J4/M, J5/S}
J1/S	{J1/S, J1/L, J1/M}
J2/S	{J2/S, J2/M, J5/S}
J3/S	{J3/S}
J4/S	{J4/S}
J5/S	{J5/S, J2/S, J5/M}

Exhibit 6

Credibility weights

Cells	Classes
J1/L	{J1/L .606, J1/M .357, J1/S .037}
J2/L	{J2/L, 1.00}
J4/L	{J4/L .278, J4/M .722}
J5/L	{J5/L, 1.000}
J1/M	{J1/M .357, J1/L .606, J1/S .037}
J2/M	{J2/M .803, J2S .197}
J3/M	{J3/M, 1.000}
J4/M	{J4/M .381, J4/L .147, J5M .472}
J5/M	{J5/M .517, J4/M .417, J5/S .065}
J1/S	{J1/S .037, J1/L .606, J1/M .357}
J2/S	{J2/S .168, J2/M .683, J5/S .150}
J3/S	{J3/S, 1.000}
J4/S	{J4/S, 1.000}
J5/S	{J5/S .112, J2/S .112, J5/M .789}