



A risk averse approach to the capacity allocation problem in the airline cargo industry

Masato Wada², Felipe Delgado^{2*} and Bernardo K. Pagnoncelli¹

¹*Escuela de Negocios, Universidad Adolfo Ibáñez, Diagonal las Torres 2640 Peñalolén, oficina 533 C, Santiago, Chile; and* ² *Departamento de Ingeniería de Transporte y Logística, Pontificia Universidad Católica de Chile, Vicuña Mackenna 4860, Macul, Casilla 306, Correo 22, Santiago, Chile*

In air cargo transportation, capacity can be reserved via allotment, which are long-term contracts with fixed price, and free, which is the space not assigned to allotment contracts. In this later case, reservations are made closer to the departure date, and normally higher tariffs are charged. The demand, the tariff, and the show-up rate for the free mode are stochastic. We consider risk neutral and risk averse formulations, using the Conditional Value-at-Risk as a risk measure. We solve the resulting problems using the Sample Average Approximation and test our models with nine experiments representing different demand patterns using real data from a major airline.

Journal of the Operational Research Society (2017) **68**(6), 643–651. doi:10.1057/s41274-016-0135-x;

published online 1 December 2016

Keywords: air transport; stochastic programming; risk

1. Introduction

The average annual growth of passengers in the airline industry in the last years has been around 5% (International Air Transport, 2012). Despite this growth, costs such as fuel have increased significantly, which yielded a net profit per passenger transported of only \$2.56 USD in 2012 (International Air Transport, 2013). Such modest results have forced airlines to search for other sources of revenue, focusing on businesses such as cargo transportation. One of the main tools that have been used to increase profits in cargo transportation is Revenue Management. However, the implementation of such tools in this industry is still in its infancy given its higher complexity when compared to passenger transportation (Kasilingam, 1997).

One important difference between the industries is the possibility that cargo customers can place reservations in two forms: allotment and free. In the allotment mode, the customer signs a long-term agreement (approximately 6 months long) specifying the weight of the load to be transported, for instance on a weekly basis, paying a fixed tariff. The free mode corresponds to the space available in the airplane that was not assigned to allotment contracts, and reservations usually take place closer to the departure date. Free reservations can, for example, be urgent deliveries, and typically airlines charge more for those types of reservations.

The decision of how much space to assign to each reservation mode is a very difficult one mainly because of

show-up rates. American Airlines, in the context of passengers, estimated that near 50% of all reservations made became no-shows or were cancelled (Smith *et al.*, 1992). In cargo, show-up rates are defined as the percentage of load that shows up in the day of the flight to be transported. Unlike passengers, for whom the show up is a binary event, in the cargo industry the show-up rate has a continuous behaviour. The actual load that the customer drops at the airport can be anything between nothing (a complete no-show) until a load greater than what was previously established (Becker and Dill, 2007). Such events happen very often in the free mode, especially because there are no penalties for any of the sides for doing so: it may happen that the airline will not transport 100% of previously agreed load, and it may be the case that the customer shows up a fraction of the load that was negotiated. The system works mainly based on trust, so the incentives for the airline to provide a good service, or for the customer to have a show-up rate of 100%, are mainly dictated by the necessity of being in good terms for future deals. To illustrate this point, suppose that, during the allotment contract, a high demand season takes place followed by a low demand one. During the first period, airlines have the incentives to accept only free bookings and not to commit themselves to allotments contracts already signed. However, if the airline does not honour the contracts, customers will send their cargo through another carrier during the low demand period, taking advantage of the excess of supply from different carriers in the market.

The impact of uncertainty can be significant to the airline: since free reservations are typically placed a few hours before the flight, the airplane may fly with empty spaces, which in a tight margin industry means losses. Therefore, airlines must aim

*Correspondence: Felipe Delgado, Departamento de Ingeniería de Transporte y Logística, Pontificia Universidad Católica de Chile, Vicuña Mackenna 4860, Macul, Casilla 306, Correo 22, Santiago, Chile.
E-mail: fdb@ing.puc.cl

at maximizing their profits by finding the right balance between the higher profit per kilogram offered by the free reservations, and the safer but less profitable allotment contracts.

There are several papers in the air cargo literature devoted to the determination of optimal overbooking levels. In the works Becker and Wald (2008) and Popescu *et al* (2006), the authors focus on the prediction of the show-up rate in order to compute the optimal overbooking levels. The publications Yunmiao Gui *et al* (2008) and Luo *et al* (2009) solve an optimization problem considering fixed capacity for weight and volume in a single-leg flight, while Wang and Kao (2008) incorporates stochastic capacity and uses fuzzy logic to determine the optimal overbooking level using scenarios to represent uncertainty. Unlike their predecessors, Zou *et al* (2013) consider multiple leg routes and solve the resulting optimization problem using a newsvendor problem with two locations. In Lei *et al* (2009), the authors propose a Conditional Value-at-Risk (CVaR) framework in which they consider spoilage and oversale costs. To the best of our knowledge, this is the only work that incorporated risk measures in the context of air cargo overbooking.

The focus of this work is to develop a planning model that provides decision makers with guidelines regarding the total space to be assigned to allotment and free reservations. We consider three sources of uncertainty: the demand, the show-up rates, and the amount charged for the free load. In the literature, some works isolate the free demand problem, and others consider allotment and free jointly. In Amaruchkul *et al* (2007), the authors use Markov decision processes to maximize the expected income of the free reservations in single-leg flights. The results show a 2.7% increase in the profits in comparison with the first-come first-booked policy. The work Huang and Chang (2010) develops a new heuristic using sampling methods, and improved the results obtained in Amaruchkul *et al* (2007). A dynamic programming model is proposed in Huang and Hsu (2005) to solve the problem under uncertain capacity. The goal is to find an optimal capacity control policy that tells when to accept an order. Their numerical results indicate that the reservations of smaller sizes should be accepted when uncertainty on demand increases. In Levin *et al* (2012), the authors develop a dynamic programming formulation that simultaneously selects allotment contracts and finds a booking control policy for the free market. Bing and Bhatnagar (2013) also use dynamic programming for the capacity booking problem, from the perspective of a freight forwarder who needs to determine his purchase strategy. The authors propose a single flight model for long-term capacity planning, and a multi flight model for the short term when more accurate demand information is revealed.

A very important point that was not considered in previous publications is the management of the risk associated with losses. The expected value is the most common tool to deal with uncertainty, but its use implies that the decision maker is risk neutral. In reality, decision makers tend to be risk averse (Lei *et al*, 2009). As pointed out by Hubbard (2009),

maximizing the expected revenue is not the only concern—it is also necessary to control the impact or the variability of losses. This is especially the case in the air cargo industry, which strongly depends on economic cycles (April IATA Economic Briefing, 2009). In the context of finance, the seminal work (Markowitz, 1952) was the first to explicitly incorporate the need to balance risk and return, using the variance as way to control risk. Another risk measure that is widely used in practice is the Value-at-Risk (VaR) (Jorion, 1997; Duffie and Pan, 1997), which is based on the quantile of the loss distribution. Although it is a very intuitive risk measure, the VaR lacks convexity, which complicates significantly its practical use.

Artzner *et al* (1999) proposed the concept of coherent risk measures, which are a set of desirable properties a risk measure should have. The Conditional Value-at-Risk (CVaR), introduced by Rockafellar and Uryasev (2000), is a coherent risk measure and therefore is suitable for optimization problems. In this work, our objective function consists of a weighted sum of the expected value and the CVaR.

Our contribution is twofold: first we propose a two-stage stochastic programming planning model to determine the total weight that should be assigned to allotment and free reservations, considering as random elements the demand, the tariff, and the show-up rate for the free mode. The model aims at giving guidelines in a planning stage about the space allocation to each type of reservation mode. Second, we depart from the classical expected value case and consider risk aversion in the objective function using the CVaR. To the best of our knowledge, our work is the first to address risk in the problem of reserving allotment space in air cargo transportation. We show that more stable solutions, with low probability of experiencing extreme losses, can be obtained at the expense of reducing the gains. We construct realistic experiments, based on real data from a major cargo company, considering nine different combinations of demand for the free mode. Each experiment reflects the presence of possible seasonalities experienced by the company, as well as situations with higher than usual variability.

The rest of the paper is organized as follows. In Section 2, we describe the problem we are modelling, discussing the assumptions we make, and a description of the data we have from the industrial partner. In Section 3, we formulate the problem and discuss how it can be solved. Section 4 presents two metrics to evaluate solution quality, and then extensive numerical experiments are performed, based on real data, for the risk neutral and risk averse cases. Finally, we make a comparison between the two models. Section 5 concludes the paper.

2. Model description

We describe the cargo problem we want to model, including the variables, the data provided by the cargo company, and how uncertainty was modelled.

2.1. Allotment and free demand

Our model assumes the cargo company is planning its cargo policy for a fixed time period in the future, say 6 months, and needs to define *before* the period starts how much kilograms will be assigned for contracts, the so-called allotment. The remaining capacity can be filled with the free demand; that is, demand that is random and is not attached to any type of contract. The price paid for the allotment is usually lower than that of the free demand, but it offers a safer alternative because the latter exhibits great variability in terms of tariffs and of show-up rates.

Airlines have different types of allotment agreements (Hellermann, 2006). In this work, we consider that allotment contracts have a fixed duration, specified at the moment the contract is signed, and they are valid for all flights that connect a specific origin and destination pair in a specific day and time within the contract's time window (Gupta, 2008). The cargo company has to decide how much of the allotment demand it will sell via contracts. In our model, we will assume there is only one contract during the time horizon under study. One can think this contract is a proxy for all contracts that were signed for this period. Thus, the allotment contract space is composed by different individual allotment contracts from different customers. During the contract's time window, some customers will arrive, others will leave, but the airline wants to use the initial estimation as a guideline for the total space dedicated to allotment.

According to the data provided, the demand magnitude, tariffs, and show-up rate for the allotment contracts do not vary significantly and are considered deterministic parameters in our model. For the show-up rates, we assume that all the assigned allotment show up for the flight, reflecting the fact that clients do not have the incentive to show up partially. If the allotment contract is not honoured by the client then the airline will end the contract or will not renew it for the following period.

On the contrary, the free demand, which is the sum of the individual booking requests, is characterized by a magnitude, a tariff and a show-up rate. The free demand is considered to be random and have different characteristics for each flight. The reason to consider an aggregated free demand is because often airlines do not keep the information needed to estimate the distribution of each booking request. Given that free demand arrives only hours before the flight, there is a significant variability in the types of reservations that arrive, e.g., urgent deliveries willing to pay higher than usual tariffs. Moreover, unlike allotment contracts, in the free reservations, partial show-up rates are common and difficult to predict, especially because there are no penalties for showing up with a fraction of the load that was reserved. Figure 1 illustrates the problem of allotment and free reservations when the number of flights in the time horizon under consideration is three. In the rectangle, representing each flight we have the demand, tariff, and show-up rate for allotment (D_A, T_A and SUR_A) and free




Figure 1 An example with three flights.

demand (D_F^i, T_F^i and SUR_F^i), $i = 1, 2$ and 3. In the next two sections, we describe how we model the random elements of the free demand.

2.2. Uncertainty modelling

The uncertain parameters of the model are all related to the free customers: the demand at each flight, the tariff that will be charged to those customers, and their show-up rate. We assume those variables to be independent, and demand and tariff are modelled as a lognormal distributions (Pak and Dekker, 2004). The use of a lognormal distribution to model individual booking requests has been widely used in the literature (Boonekamp *et al.*, 2013; Levin *et al.*, 2012; Pak and Dekker, 2004). Following Cobb *et al.* (2012) we assume that the free demand of a flight, which is the sum of the individual booking requests, follows this same distribution with the same first and second moments. Aside from being positive and allowing values that are significantly greater than the mean, the lognormal distribution generates values that resembles the data we encountered in practice. The same applies to the tariff charged to free customers.

Finally, to model the *show-up rate*, we use a discrete distribution similar to the one proposed by Popescu *et al.* (2006). To generate this distribution, we create different bins, representing show-up rate ranges, with different probabilities of occurrence to each one. To set the optimal number of bins, we started with the methodology developed by Birgé and Rozenholc (2006) that uses bins of equal size and then adjusted bin sizes due to the scarcity of data available.

2.3. Data description

The test data for this model are based on real data provided by a major cargo airline for a single market between December 2013 and March 2014. In this market, three flights are scheduled each week. In total, we collected 10011 shipment orders of the 37 flights flown during this period. The database includes information of each flight, with specific details of each shipment order: the client name, number of the order,

flight date, shipment origin and destination, load reserved, load flown, tariff, and reservation mode. However, the database does not record the date of each booking, and bookings that were not accepted. For confidentiality reasons, we used synthetic data based on the information received from the airline.

For all our experiments, we consider that the flight capacity (C) is fixed and equals to 100000 kg, representing the capacity of a CAO Boeing 777, while V represents the number of flights under consideration. We set $V = 3$ emulating the presence of three different seasons of equal length during the planning process where the allotment contracts take place. Our commercial partner plans their scheduling in seasons, representing events such as holidays, harvesting. Although the model can handle larger number of seasons, $V = 3$ captures the variability observed in the data we had access. Table 1 summarizes the values of the deterministic parameters.

Regarding free reservations, we consider that $D_F \sim \ln(11.32 \text{ (kg)}, 0.365^2)$ and $T_F \sim \ln(1.525 \text{ (USD/kg)}, 0.044^2)$, where $\ln(\mu, \sigma^2)$ represents a lognormal random variable with parameters μ and σ estimated based on the data we had available. The show-up rate follows a discrete distribution as was explained in Section 2.2. Table 2 shows the show-up rate ranges with the probabilities of occurrence. From the table, we can observe that the last bin allows $\text{SUR}_F > 1$ reflecting the possibility that certain customers arrive with more load than reserved. To obtain one sample for the random variable show-up rate, two independent random numbers are used. The first follows a $U(0, 1)$ and is used to identify the bin, and the second is uniform between the lower and upper bounds that define the value.

3. Formulation

In this section, we show the explicit optimization formulation of the problem, and discuss how the Sample Average Approximation will be used to solve it.

3.1. Risk neutral formulation

The risk neutral formulation aims at maximizing the average income obtained with contracts plus free demand. The problem fits into the framework of two-stage stochastic programming: the airline has to decide how much should be assigned for allotment *before* uncertainty is revealed, and when the values of the random elements are known, the assignment for the free demand must be chosen. We formulate the problem as follows:

$$\text{Max}_{0 \leq X_A \leq D_A} T_A X_A \text{SUR}_A + \frac{1}{V} \sum_{i=1}^V \mathbb{E}[Q(X_A, \omega_i)] \quad (1)$$

The constraint means that the capacity X_A destined for allotment contracts must be smaller than or equal to the total demand D_A for such contracts. The objective function of (1) has two terms: the first represents the profit obtained through contracts in any flight, which is the tariff T_A multiplied by the assigned capacity times the show-up rate for allotment SUR_A , and the second represents the average return per flight obtained through free contracts. Thus, the objective function value can be interpreted as the average income per flight when using a contract with demand D_A and tariff T_A . Function $Q(X_A, \omega_i)$ represents the profit obtained with one flight *given* amount reserved for the contract X_A , and realizations of demand, tariff, and show-up rate for the free customers for each flight ω_i . The second stage can be explicitly written as the optimal value of a maximization problem:

$$Q(X_A, \omega_i) = \text{Max}_{0 \leq X_F^i \leq D_F^i} T_F^i X_F^i \text{SUR}_F^i \quad (2)$$

$$\text{s.t. } X_A \text{SUR}_A + X_F^i \text{SUR}_F^i \leq C. \quad (3)$$

Constraint (3) states that the sum of effective cargo for allotment and free customers cannot exceed the airplane's capacity C .

The random variables are continuous, and the expectation of the function $Q(X_A, \cdot)$ is impossible to compute explicitly. Following Linderoth *et al* (2006), we can approximate the problem using the Sample Average Approximation (SAA) method. Each sample j represents a vector of three components: $(D_F^{ij}, \text{SUR}_F^{ij}, T_F^{ij})$. For N samples, the problem becomes

$$\begin{aligned} & \text{Max}_{0 \leq X_A \leq D_A, X_F^{ij}} T_A X_A \text{SUR}_A + \frac{1}{V} \frac{1}{N} \sum_{i=1}^V \sum_{j=1}^N T_F^{ij} X_F^{ij} \text{SUR}_F^{ij} \\ & \text{s.t. } X_A \text{SUR}_A + X_F^{ij} \text{SUR}_F^{ij} \leq C, \quad i = 1, \dots, V, \quad j = 1, \dots, N, \\ & 0 \leq X_F^{ij} \leq D_F^{ij}, \quad i = 1, \dots, V, \quad j = 1, \dots, N. \end{aligned}$$

The expected value model is a completely valid formulation, but as discussed before it does not take into account risk. We now turn our attention to the risk averse model with CVaR.

3.2. Risk averse formulation: the λ -CVaR model

The λ -CVaR model incorporates risk in the objective function, using the weight λ to balance risk and return. For the sake of completeness, let us define the CVaR rigorously. The Value-

Table 1 Deterministic parameters

Parameter	D_A	T_A	SUR_A	C	V
Value	51847 (kg)	2.5 (USD/kg)	1	100000 (kg)	3

Table 2 SUR_F distribution

Range	[0.46, 0.65)	[0.65, 0.80)	[0.80, 0.90)	[0.90, 1.00)	[1.00, 1.08]
Probability	0.14	0.31	0.19	0.31	0.05

at-Risk of a random variable X that represents losses at confidence level α is defined as:

$$\text{VaR}_\alpha[X] := \min\{t \mid F(t) \geq \alpha\} = \min\{t \mid P(X \leq t) \geq \alpha\}. \quad (4)$$

The Conditional Value-at-Risk (CVaR) of a continuous random variable X with cdf $F(\cdot)$, and risk level $\alpha \in [0, 1]$, is defined as the average of losses given that such losses were higher than the VaR:

$$\text{CVaR}_\alpha[X] = \mathbb{E}[X \mid X > \text{VaR}_\alpha[X]]. \quad (5)$$

Albeit simple, definition (5) is of little use in an optimization context. The key result in Rockafellar and Uryasev (2000) is the proof that CVaR can be expressed as the optimal value of the following optimization problem:

$$\text{CVaR}_\alpha[X] = \min_{\theta \in \mathbb{R}} \left\{ \theta + \frac{1}{1-\alpha} \mathbb{E}[(X - \theta)_+] \right\}, \quad (6)$$

where $(a)_+ := \max(a, 0)$. Since we assume positive values are losses, we switch to minimization and write the λ -CVaR model as

$$\begin{aligned} \min_{0 \leq X_A \leq D_A} & -T_A X_A \text{SUR}_A + \frac{1}{V} \sum_{i=1}^V \lambda \mathbb{E}[-Q(X_A, \omega_i)] + (1 \\ & - \lambda) \text{CVaR}_\alpha[-Q(X_A, \omega_i)]. \end{aligned} \quad (7)$$

Using expression (6) and defining auxiliary variable u_{ij} to linearize the CVaR, problem (7) can be approximated by

$$\begin{aligned} \min_{X_A, X_F^{ij}} & -T_A X_A \text{SUR}_A + \frac{1}{V} \sum_{i=1}^V \left\{ \lambda \left[\frac{1}{N} \sum_{j=1}^N -T_F^{ij} X_F^{ij} \text{SUR}_F^{ij} \right] + (1 - \lambda) \right. \\ & \left. \left[\theta + \frac{1}{N(1-\alpha)} \sum_{j=1}^N u^{ij} \right] \right\} \\ \text{s.t. } & 0 \leq X_A \leq D_A, \\ & u^{ij} \geq 0, \quad i = 1, \dots, V, \quad j = 1, \dots, N, \\ & u^{ij} + \theta + T_F^{ij} X_F^{ij} \text{SUR}_F^{ij} \geq 0, \quad i = 1, \dots, V, \quad j = 1, \dots, N, \\ & X_A \text{SUR}_A + X_F^{ij} \text{SUR}_F^{ij} \leq C, \quad i = 1, \dots, V, \quad j = 1, \dots, N, \\ & 0 \leq X_F^{ij} \leq D_F^{ij}, \quad i = 1, \dots, V, \quad j = 1, \dots, N. \end{aligned}$$

Such formulation can be efficiently solved by standard optimization software, and different values of the parameter λ generate different risk profiles. The choice of λ is problem dependent, and there are very few studies on the topic. In Shapiro *et al.* (2015), the authors run several numerical experiments for a hydrothermal scheduling problem and propose a criteria to find the optimal choice of the parameter

based on some problem-dependent regret criteria. They also note that for values too close to one the problem is essentially risk neutral, and for choices close to zero risk averseness increases costs significantly, generating solutions with little practical value. The numerical results with different values of λ will be shown in Section 4.

4. Numerical experiments

4.1. Benchmark policies

In the literature of stochastic programming, it is common to use two metrics to obtain additional information about the usefulness of a model. The first metric is the Value of the Stochastic Solution (VSS), which is a measure of the gain in considering randomness in the formulation. If the VSS is zero then replacing the random parameters by their averages yields the same results as the two-stage stochastic programming model.

The other metric is the Expected Value of Perfect Information, or EVPI. The EVPI measures how much a decision maker would be willing to pay for perfect information, that is, for accessing the outcomes of the random variables before they are revealed. An EVPI of zero means that the solution of the two-stage problem, which is obtained *without* knowing the outcomes of the random variable, is as good as the solution with complete information. More details can be found in Birge and Louveaux (2011). In the next section, we will present our numerical results, including the value of both metrics for each of our experiments.

4.2. Experiment design

In order to test the model under a broader array of demand patterns, we construct a set of nine different experiments. Each of the nine experiments consists in choosing one demand category for each of the three flights. The categories, which are expressed in Table 3, represent different demand patterns for the free demand, depending on the mean and standard deviation of the underlying random variable.

In the base category MM, we assume a demand mean of 88560 kg and a standard deviation of 33503 kg. Both values were obtained from the data provided by our industrial partner. To generate the high (low) mean demand level, we took the base demand level and amplified (reduced) by 25%. Analogously, high (low) levels of standard deviation were created by increasing (decreasing) the coefficient of variation by 15% with respect to the base case level.

Observing patterns in the data, we built nine experiments, displayed in Table 4. They were chosen to represent specific demand patterns as a result of events such as Mother's Day, San Valentine's day or harvesting season, that generate peaks and valleys on demand. For instance experiment 2 the mean demand is equal to the base case, but the standard deviation is high for all three flights, while in experiment 5 we have a changing pattern for the mean. We will test risk neutral and risk averse models presented in Section 3 for all nine experiments described in Table 4. The results of these experiments are presented in the next section.

4.3. Results

As is usually the case with stochastic programming problems, we cannot solve the resulting problem directly since the expected value cannot be computed explicitly. As mentioned in Section 3.1, we will approximate the problem using SAA in order to obtain good candidate solutions, as well as statistical upper bounds for the true optimal value. We use the risk neutral formulation and experiment 1 to test the influence of the number of samples N and the number of experiments M in the stability of the objective function. To solve the optimization problem, we use Gurobi 5.6.2 (Gurobi Optimization, 2012).

4.3.1. Risk neutral results We use SAA to obtain lower and upper bounds for the unknown optimal value v^* of problem (1) for each of the nine experiments. For the upper bound, we ran $M = 100$ experiments with $N = 500$, which generated 100 candidate solutions and 100 objective function values. The values of M and N were chosen for computational tractability. We performed experiments with different values of M and N ,

and the results were similar to the ones we are reporting. According to Linderoth *et al* (2006), the average of those values is an (statistical) upper bound for v^* , and confidence intervals can be constructed. Among those 100 candidates generated, we selected the best candidate \hat{x}_A^* in terms of objective function value, which is feasible by construction, to compute the lower bound. By fixing the allotment capacity equals to \hat{x}_A^* , we performed out-of-sample experiments with much higher value of $N' > N$ to construct accurate statistical estimates of the objective function value associated with \hat{x}_A^* , including confidence intervals.

By constructing lower and upper bounds, we have not only a candidate solution that can be implemented in practice, but we also can infer about the quality of such candidate by comparing its objective function value with the upper bound. Results for each of the nine experiments considering a value of $N' = 1 \times 10^6$ are shown in Table 5.

In the second column, $f(\hat{x}_N^j)$ represents the objective function value evaluated at candidate point \hat{x}_N^j . We report the best candidate obtained over M runs, for each case. The third column is the statistical upper bound, obtained as described in Linderoth *et al* (2006). The last column uses information of the first two to obtain a statistical optimality gap for each experiment.

From Table 5, we observe that the length of the confidence intervals for the upper bounds is around ten times larger than the ones for the objective function value estimation. The reason is due to the difference in the values of sample sizes N and N' . Note that in the former case, we have to solve optimization problems, while in the latter we are simply generating a statistical estimate. Tighter estimates can be obtained using variance reduction techniques such as Latin Hypercube Sampling (Linderoth *et al*, 2006). More

Table 3 Free demand categories

Demand mean	Standard deviation		
	High (H)	Medium (M)	Low (L)
High (H)	HH	HM	HL
Medium (M)	MH	MM	ML
Low (L)	LH	LM	LL

Table 4 Description of free demand experiments

Experiment	Characteristic	Flight 1	Flight 2	Flight 3
1	Base case	MM	MM	MM
2	Variability increase	MH	MH	MH
3	Demand increase	HM	HM	HM
4	Demand decrease	LM	LM	LM
5	3 different seasons	MM	LM	HM
6	1 high demand season	MM	HM	MM
7	1 low demand season	MM	LM	MM
8	3 seasons, high variability	MH	LH	HH
9	Variability decrease	ML	ML	ML

importantly, lower and upper bounds for each of the nine experiments are remarkably close, showing that SAA provides excellent approximations for this problem. The (stochastic) optimality gap as a percentage of the approximate objective function value $f(\hat{x}_N^j)$ is never greater than 0.5% for all nine experiments.

Table 6 shows the results compared with our benchmark policies explained in Section 4.1. The *EVPI* in this case represents a sample estimation of the maximum willingness to pay for having perfect information of the demand magnitude, tariffs, and show-up rates that a particular flight will face. The results shows us that this willingness to pay considering all nine experiments is on average 37,690 [USD] per flight. In contrast, the sample estimate of the VSS yields on average a 1% gain for considering the randomness in the problem. Nevertheless, the VSS increases in scenarios 2 and 8, in which the free demand has higher variability, and also present seasonalities, reaching a 1.64% and 1.35% improvement.

4.3.2. Risk averse results In this section, we discuss the influence of considering risk in the objective function using the CVaR. First, we study the influence of the weight λ , which gives the desired balance between risk and return. Second, we explore the impact of the confidence level α , which calibrates how strict the CVaR measure is with respect to extreme losses. For both analysis, we use experiment 1, which is the base case experiment. Our main goal is to understand how the inclusion of risk affects the first stage decision, that is, the percentage destined to the allotment.

Figure 2 shows how the optimal allotment capacity changes for different values of λ , considering a confidence level $\alpha = 0.95$. Note that a value of $\lambda = 1$ represents the risk neutral approach where only returns matter, whereas $\lambda = 0$ only gives importance to risk. For $\lambda \leq 0.5$, which represents high levels of risk aversion, the optimal solution is to assign the maximum available capacity to allotment. The intuition behind this result is clear: risk averse decision makers try to avoid variability, and the allotment is a stable, albeit less profitable, strategy. As the parameter λ increases the weight of the CVaR decreases, and the allotment allocation is much smaller. Less risk averse decisions makers are willing to bet

that the more profitable free demand will generate greater profits. As discussed in Shapiro *et al* (2015), values of λ close to one or to zero are probably too extreme to be used in practice.

The value of α indicates that the average of losses greater or equal than the α -quantile will be considered. Values closer to one imply that the decision maker is extremely risk averse and wants to be protected against extreme losses. In the limit, the CVaR turns into the worst-case risk measure. Figure 3 shows the allotment capacity that should be reserved for different values of α , considering a fixed $\lambda = 0.6$. It is interesting to see the sharp increase towards the maximum allotment capacity as α approaches one. Even for $\lambda = 0.6$ we observe that a small value of α , e.g., 0.3, leads to an aggressive allotment allocation around 40%.

4.3.3. Comparison between risk neutral and risk averse results So far, we have not analysed the optimal value of the λ -CVaR model because, except in the case where $\lambda = 1$ where the problem is identical to the risk neutral model, the objective function value of the λ -CVaR model has no physical, i.e., monetary interpretation. It is fundamental to have an estimation of the expected gain for the CVaR solution in order to observe the risk return trade-off of the model. We showed




Figure 2 Allotment capacity as a function of λ for $\alpha = 0.95$.




Figure 3 Allotment capacity as a function of α for $\lambda = 0.6$.

Table 5 Lower and upper bound estimates for v^*

Experiment	Best $f(\hat{x}_N^j)$ (95% CI)	Upper bound (95% CI)	Optimality gap (%) (95% CI)
1	353779 \pm 57	354360 \pm 443	0.30
2	339820 \pm 62	340490 \pm 496	0.36
3	379334 \pm 75	380491 \pm 608	0.48
4	328087 \pm 41	328160 \pm 348	0.14
5	347937 \pm 51	348392 \pm 397	0.25
6	360925 \pm 63	361395 \pm 530	0.29
7	343086 \pm 56	343286 \pm 415	0.19
8	335351 \pm 72	335538 \pm 494	0.22
9	368109 \pm 53	368773 \pm 408	0.30

Table 6 Results for the risk neutral formulation

<i>Experiment</i>	<i>EVPI</i>	<i>VSS</i>
1	39137	2434
2	41323	4608
3	38580	3062
4	26052	1828
5	40613	3191
6	40734	2905
7	37259	2533
8	41545	5489
9	33968	733
Average	37690	2975

Table 7 Comparison results between risk neutral and risk averse model

<i>Experiment</i>		<i>Dif RN (%)</i>	<i>Dif $\lambda = 0.7$ (%)</i>
1	Income	-0.57	2.01
	SD	-20.86	-57.33
2	Income	-1.60	-0.02
	SD	-27.74	-56.95
3	Income	-0.65	2.37
	SD	-20.66	-58.78
4	Income	-0.24	-0.12
	SD	-22.64	-29.14
5	Income	-1.03	1.25
	SD	-23.10	-59.14
6	Income	-1.21	0.83
	SD	-22.97	-62.34
7	Income	-0.91	0.18
	SD	-23.55	-55.78
8	Income	-2.28	-1.41
	SD	-29.68	-53.87
9	Income	-0.14	1.95
	SD	-12.45	-51.09

that the CVaR offers protection against variability; now we want to estimate the cost of such protection.

We fix the allotment decision $x_A^{RA} = 49.2\%$ obtained for $\lambda = 0.7$ and the allotment decision $x_A^{RN} = 33\%$ obtained for the risk neutral case. We then generate 100 new experiments, each with 500 scenarios, and compute the expected income obtained by those two solutions. Finally, we compute the average income and standard deviation over all scenarios. The results are displayed in Table 7.

Table 7 shows for each experiment the percentage difference for both the expected income per flight and the standard deviation of this income compared against the solution that ignores uncertainty and optimizes over the average values of the parameters. We will refer to this solution as the benchmark case for this analysis. A negative percentage value in the income in Table 7 represents greater profits than the benchmark solution, while a negative percentage in the standard deviation means a reduction with respect to the benchmark case.

In terms of income, the expected value solution is on average 1% greater than the benchmark solution. In the case of the expected value the greater benefits occur for experiments with high standard deviation and stationarity, while lower benefits occurred for experiments 1 and 9, for which the standard deviation is moderate or low, and there is no stationarity. The average reduction in terms of standard deviation is 22.7%, which shows that, by ignoring uncertainty, the benchmark case faces a much higher variability of gains.

For the CVaR case, even though on average the income is 0.8% worse, we observe that when facing high variability and high stationarity, such as in experiments 2 and 8, the solution is better than the benchmark case. More striking is the reduction in the variability obtained by the CVaR solution, which is on average 53.8% better than the benchmark case. In summary, even though the income of the CVaR case can be smaller, the decision maker adopting this strategy will receive a significantly more stable cash flow. This is particularly important in air cargo industry which, as mention before, is very susceptible to economic cycles.

5. Conclusions

In this paper we study the air cargo transportation problem from a tactical perspective and propose a two-stage risk averse stochastic programming model to determine the total weight that should be assigned to allotment and free reservations over a given time horizon. We depart from the classical expected value formulation and consider risk aversion in the objective function, by means of the Conditional Value-at-Risk (CVaR). We construct nine experiments under different demand patterns, including stationarity, and solve the resulting problems using the Sample Average Approximation (SAA) technique. We present numerical results for each of the nine experiments, including lower and upper bounds for the optimal value of those problems. We also compare the results against a decision maker that ignores uncertainty and replaces the uncertain parameters by its averages.

We show that SAA is a powerful technique to deal with the air cargo transportation problem: the (stochastic) gap obtained for the risk neutral formulation was never greater than 0.5% of the objective function value of the candidate solutions we found for any of the nine experiments. In addition, we compute the Expected Value of Perfect Information (EVPI) and the Value of the Stochastic Solution (VSS) for all cases. The EVPI was around 10% of the optimal values of the stochastic problems, and we show that there is a 1% gain on average in considering randomness. The value increases for experiments with higher variability.

For the risk averse model, we study the behaviour of the solution as we change the risk aversion parameter λ , and the

CVaR risk level α . For smaller values of λ , the solution reserves a high percentage of allotment, which yields safer and more stable incomes. As λ approaches one and the weight of the CVaR in the objective function decreases, the percentage of allotment decreases significant, converging to the value obtained in the risk neutral solution. Similar results were obtained for the risk level α .

Finally, we performed a more detailed comparison concerning the expected income and the variability of those incomes for the risk averse solution compared to the fully deterministic case, in which all random parameters are replaced by their average values. Unsurprisingly, the risk averse model generates smaller incomes on average, which is essentially the price for being protected against risk. More relevant is the reduction in the variability of incomes: by performing 100 experiments, the standard deviation of returns is 53.8% smaller on average. For a decision maker wishing a stable flow of returns, the CVaR solution is clearly a desirable alternative, especially on markets that face high variability and stationarity.

References

- Amaruchkul K, Cooper W, and Gupta D (2007). Single-leg air-cargo revenue management. *Transportation Science* **41**(4):457–469.
- Artzner P, Delbaen F, Eber J-M and Heath D (1999). Coherent measures of risk. *Mathematical Finance* **9**(3):203–228.
- Becker B and Dill N (2007). Managing the complexity of air cargo revenue management. *Journal of Revenue & Pricing Management* **6**(3):175–187.
- Becker B and Wald A (2008). Air cargo overbooking based on the shipment information record. *Journal of Revenue & Pricing Management* **7**(3):242–255.
- Bing L, Bhatnagar R (2013). Optimal capacity booking in air cargo transportation. *Journal of Revenue & Pricing Management* **12**(3):235–253.
- Birge J and Louveaux F (2011). *Introduction to Stochastic Programming*. Springer Science & Business Media.
- Birgé L and Rozenholc Y (2006). How many bins should be put in a regular histogram. *ESAIM: Probability and Statistics* **10**:24–45.
- Boonekamp T, Gromicho J, Dullaert W, and Radstaak B (2013). *Air Cargo Revenue Management*. Unpublished masters thesis, Vrije Universiteit, Amsterdam.
- Cobb B, Rumi R and Salmerón A (2012). Approximating the distribution of a sum of log-normal random variables. *Statistics and Computing* **16**(3):293–308.
- Duffie D and Pan J (1997). An overview of value at risk. *The Journal of Derivatives* **4**(3):7–49.
- Gui Y, Gong B and Cheng Y (2008). Research on air cargo overbooking model. In: *4th International Conference on Wireless Communications, Networking and Mobile Computing*, 2008. WiCOM'08. IEEE, pp. 1–4.
- Gupta D (2008). Flexible carrier-forwarder contracts for air cargo business. *Journal of Revenue and Pricing Management* **7**(4):341–356.
- Hellermann R (2006). *Capacity Options for Revenue Management: Theory and Applications in the Air Cargo Industry*, vol 575. Springer Science & Business Media.
- Huang K and Chang K (2010). An approximate algorithm for the two-dimensional air cargo revenue management problem. *Transportation Research Part E: Logistics and Transportation Review* **46**(3):426–435.
- Huang K and Hsu W (2005). Revenue management for air cargo space with supply uncertainty. In: *Proceedings of the Eastern Asia Society for Transportation Studies*, vol 5, pp. 570–580.
- Hubbard D (2009). *The Failure of Risk Management: Why it's Broken and How to Fix It*. John Wiley and Sons.
- April IATA Economic Briefing (2009). www.iata.org/economics. Air freight timely indicator of economic turning point.
- International Air Transport Association (2012). Air transport market analysis december 2012.
- International Air Transport Association (2013). Special report: airline profitability.
- Jorion P (1997). *Value at Risk: The New Benchmark for Controlling Market Risk*. Irwin Professional Pub.
- Kasilingam R (1997). Air cargo revenue management: characteristics and complexities. *European Journal of Operational Research* **96**(1):36–44.
- Lei L, Zhou J and Wang H (2009). Air cargo overbooking model with stochastic capacity and penalty cost under CVaR framework. In: *Proceedings of 2009 IEEE International Conference on Grey Systems and Intelligent Services*. IEEE, pp. 1234–1240.
- Levin Y, Nediak M, Topaloglu H (2012). Cargo capacity management with allotments and spot market demand. *Operations Research* **60**(2):351–365.
- Linderoth J, Shapiro A and Wright S (2006). The empirical behavior of sampling methods for stochastic programming. *Annals of Operations Research* **142**(1):215–241.
- Luo S, Çakanyıldırım M and Kasilingam R (2009). Two-dimensional cargo overbooking models. *European Journal of Operational Research* **197**(3):862–883.
- Markowitz H (1952). Portfolio selection. *The Journal of Finance* **7**(1):77–91.
- Optimization G et al (2012). Gurobi Optimizer Reference manual (2012). <http://www.gurobi.com>.
- Pak K and Dekker R (2004). Cargo Revenue Management: Bid-Prices for a 0-1 Multi Knapsack Problem. ERIM Report Series Reference No. ERS-2004-055-LIS. Available at SSRN: <https://ssrn.com/abstract=594991>.
- Popescu A, Keskinocak P, Johnson E, LaDue M and Kasilingam R (2006). Estimating air-cargo overbooking based on a discrete show-up-rate distribution. *Interfaces* **36**(3):248–258.
- Rockafellar T and Uryasev S (2000). Optimization of conditional value-at-risk. *Journal of risk* **2**(3):21–42.
- Shapiro A, Cabral F and da Costa J (2015). *Guidelines for Choosing Parameters λ and α for the Risk Averse Approach*. Report, Georgia Institute of technology and ONS.
- Smith B, Leimkuhler J and Darrow R (1992). Yield management at american airlines. *Interfaces* **22**(1):8–31.
- Wang Y-J and Kao C-S (2008). An application of a fuzzy knowledge system for air cargo overbooking under uncertain capacity. *Computers & Mathematics with Applications* **56**(10):2666–2675.
- Zou L, Yu C, and Dresner M (2013). The application of inventory transshipment modeling to air cargo revenue management. *Transportation Research Part E: Logistics and Transportation Review* **57**:27–44.

Received 11 June 2015;
accepted 4 October 2016