Modelling Aircraft Loans & Leases

Nils Hallerstrom

January 2020

Preface

These discussion notes are an update to notes that I first prepared in 2010. The issues discussed here are ones that I have pondered over the last 25 years, trying to model aircraft loans and leases quantitatively. In 1993, Jan Melgaard (then at PK) and I worked with Bo Persson of Xice AB in Sweden to develop an analytic model of aircraft loans that we called SAFE. This model evolved into a Monte Carlo simulation tool, Lending EDGE, that was taken into operation at PK in 2012 and validated under ISRS 4400 by Deloitte in 2013. After my retirement from PK AirFinance at the end of 2018, Bo Persson of Xice AB and I have subsequently developed a different but related model for structuring, pricing and optimizing aircraft backed loans and leases called MidaX.

I have made some corrections and amendments to the 2010 notes, based on helpful feed-back from industry practitioners and academics, expanded on maintenance value, and, based on Monte Carlo technology, added a section on interrelationships among aircraft values, obligor default and interest rates to permit modeling of diversification effects, cross-collateralization and portfolios.

My work at PK AirFinance taught me a lot about risks and rewards in aircraft finance, not least from the deep experience and insight of many valued customers and co-workers at PK and GECAS, PK's then parent company, but the views and opinions expressed herein are my own.

In preparing these notes, I have been helped by several people with whom I have had many inspiring discussions. I would like to mention in particular Bo Persson of Xice AB, Christophe Beaubron at GECAS, Seth Aslin (at the time with White Kite), Professor Vadim Linetsky of Northwestern University, Professor Steven Golbeck of University of Washington, Professor Alessandro Gavazza of New York University, Mats Levander of the Swedish Central Bank, and Alan Picone (at the time with Deloitte, now with KPMG). None of these individuals have reviewed, nor approved the final text, so any errors or misconceptions are entirely my own.

I would welcome any feedback, comments and ideas.

Luxembourg, January 2020

Nils Hallerstrom

nils@hallerstrom.com

Abstract

In these notes, I discuss how airline defaults, aircraft values, and interest rates can be forecasted and simulated. An aircraft capacity surplus/shortage cycle is defined and I discuss how this can be forecasted and simulated. Then we look at how this cycle may impact airline defaults, aircraft values, and maybe interest rates, and how it could cause correlation among those. We then turn to pay-off functions for aircraft loans and leases in the event of an airline default or contract maturity with an unamortized residual book value or non-recourse balloon payment. With the simulated scenarios or forecasted distributions of default probabilities, aircraft values and interest rates, knowing the pay-off function, I then show how the net present value of aircraft loans or leases can be modeled as probability distributions. From these distributions, I discuss how various measures for risk, value, and performance can be derived. I then discuss briefly how these measures could be used for pricing, structuring, underwriting, and valuing aircraft loans and leases, and how one could set reserves for expected losses. Finally, I point to an approach to analyze a portfolio of Aircraft Backed Loans, accounting for interrelationships among aircraft, obligors, and interest rates.

NB: These discussion notes are, by their nature, general and are not intended to be, nor should they be, relied upon as advice. Specific guidance and consideration should always be sought and applied to particular cases, circumstances, asset types, jurisdictions, etc.

Contents

Introduction

When evaluating aircraft loans, it is typical to consider the cash flows from the disbursement of the initial loan amount, principal repayments, interest and margin payments, fees, and calculate the net present value (NPV) or the internal rate of return (IRR). In an operating lease, we take the initial purchase price of the aircraft, rental payments, interest, maintenance reserve payments and expenses, and the sales proceeds from a sale at the assumed residual value at the end of the lease term to calculate the NPV or IRR. Taxes, refinancing terms and capital rules can also be accounted for. Based on internal targets, terms can be structured to meet criteria to accept or reject a proposed transaction.

The limitations with this approach are that 1) the contractual debt service or rental payments are sometimes interrupted when the airline **defaults**, 2) the **aircraft value** at the end of the term, or when a default occurs, is not necessarily what was assumed at the outset, and 3) the **interest rate** may change over time and thus affect the net present value. These uncertainties create risks - default risk, aircraft value risk and interest rate risk. Therefore, the NPV and the IRR are not static numbers, but ranges or probability distributions of possible outcomes.

This paper will discuss how these uncertainties in airline survival, aircraft value and interest rates could be estimated, how we can establish the NPV/IRR as a distribution, and how this distribution can be used as a guide to **structure**, **price**, **accept** or **reject** a proposed transaction, how to measure the **value** of a risky loan or lease, or how to measure **risk** for purposes of setting **loss reserves** or **allocating capital**. We first discuss stand-alone transactions with a single counterparty and a single aircraft. The issues around multiple aircraft and airline transactions, cross collateralization, portfolios, diversification, and concentrations, where the interdependence of aircraft values, airline defaults and interest rates come into play is briefly discussed at the end. We will ignore tax risks and operational risks.

Fundamentally, future default events, aircraft values and interest rates are guesses we make.

History is a guide, but we live in an ever-changing world where economics, politics, and social patterns are shaped by humans and not by the time invariant laws of physics. The history of the jet age is so short that past events cannot possibly cover all possible scenarios of the future. Experience, judgment, and instinct will have to complement the historical data in order to make meaningful assumptions about future outcomes.

Market prices of financial assets can also be a guide. Prices of Credit Default Swaps or Interest Rate Derivatives can help us make assumptions about default expectations or the evolution of interest rates. Prices of EETCs could reveal something about the market's aircraft value expectations, although default risk and interest rates will influence the price as well.

However, even with the best data and assumptions, models alone will not be enough to succeed in aircraft financing. We will point to tools and concepts that can bring us beyond the static analysis where risk and reward are only accounted for in intuitive ways, and help us make faster, better, and more coherent decisions. Such models will help us be consistent 1) **across transactions** in how we view a particular airline or aircraft type, they will help us to be consistent 2) **within transactions** in how we weigh the impact of one feature against another, credit quality, loan-to-value ratio, aircraft quality, pricing and term structure, and blend all the moving parts of a transaction into a coherent picture, and finally, they will help us be consistent 3) **over time**.

Experienced transactors can evaluate features such as pricing, Loan to Value Ratio (LTV), desirability of an aircraft type, credit worthiness of the airline, and risk mitigants such as

security deposits, first loss deficiency guarantees, or residual value guarantees etc. However, it is humanly impossible to intuitively trade one feature against another in a coherent and consistent manner. The relationships between the different features are non-linear, time-dependent, and sometimes non-intuitive. The combinations of features are unlimited. This is why a model to structure, price and evaluate transactions can be very useful. The guess-work is broken down into the constituent parts of the deal, and the model will help us structure our thinking around a complex transaction.

In the following, I will discuss how we can forecast airline defaults, future aircraft values, and interest rates, with associated probability distributions. We will look at how these properties are interdependent and what may cause this correlation. Correlation is problematic because when airlines default, aircraft values have a tendency to be low, and often interest rates are low, to a point where the interest rate swap we entered into to hedge our fixed rate carries breakage costs. We will then look at the pay-off distribution when we face default or maturity of a loan or lease with a remaining "balloon" payment or a residual book value. With this, we can then build a probability distribution of the net present value of the cash flows generated by the deal under different scenarios. And finally, I will discuss how this probability distribution can be used to measure expected return, risk and value in a transaction.

I will not prescribe any particular model, merely point to different approaches that are possible, and discuss shortcomings and opportunities.

Financial models have taken some blame for contributing to the recent financial crisis. Any model is an incomplete description of a complex reality. It is important to understand the limitations of models, but also to understand the advantages they can offer. Whenever we decide to offer a loan or a lease on specified terms, we have implicitly made assumptions about what we think about the future. It is better to make those assumptions explicitly. A key to successfully operating a model is that input assumptions are honest and consistent. This is easiest to achieve when the same basic model is used for pricing, underwriting, as for portfolio monitoring.
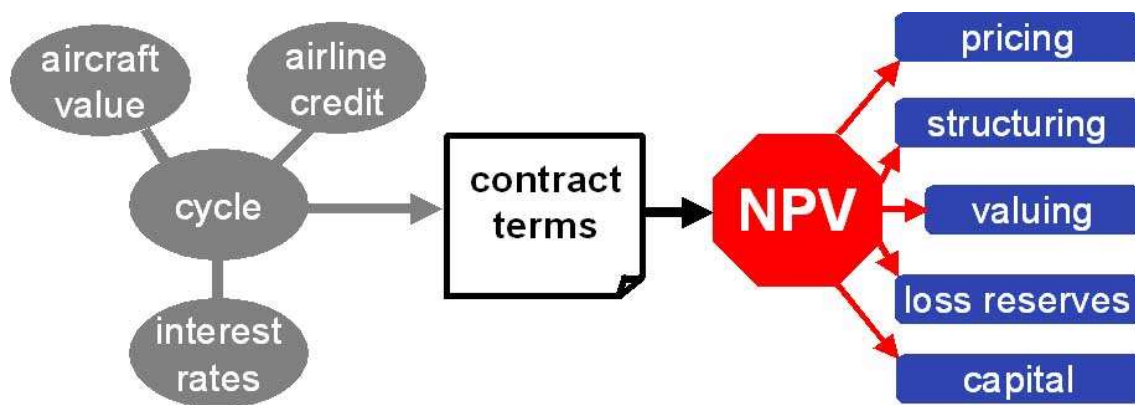
Products

Up until the 1970's, airlines would mainly finance their aircraft with equity and corporate debt. Loan to value ratios (LTVs) were around 50%, but the banks would primarily rely on the airline credit. With deregulation and increasing competition came a demand for financing to non-investment grade credits, and from this, asset-based lending and operating leasing emerged. Aircraft-backed finance can take many shapes and forms, but the idea is that the aircraft itself can be looked upon as a source of repayment, either at default or at the maturity of a financing, when the contractual obligation by the airline to pay debt service or rent has ended, but a residual amount is still outstanding.

The estimated value of the current world fleet of commercial jet aircraft is around $520bn. New deliveries are valued at around $95bn per annum, so with re-financings, it is safe to say that aircraft finance is a $100bn+ business. The picture below shows examples of providers, product groups and product sub-groups.
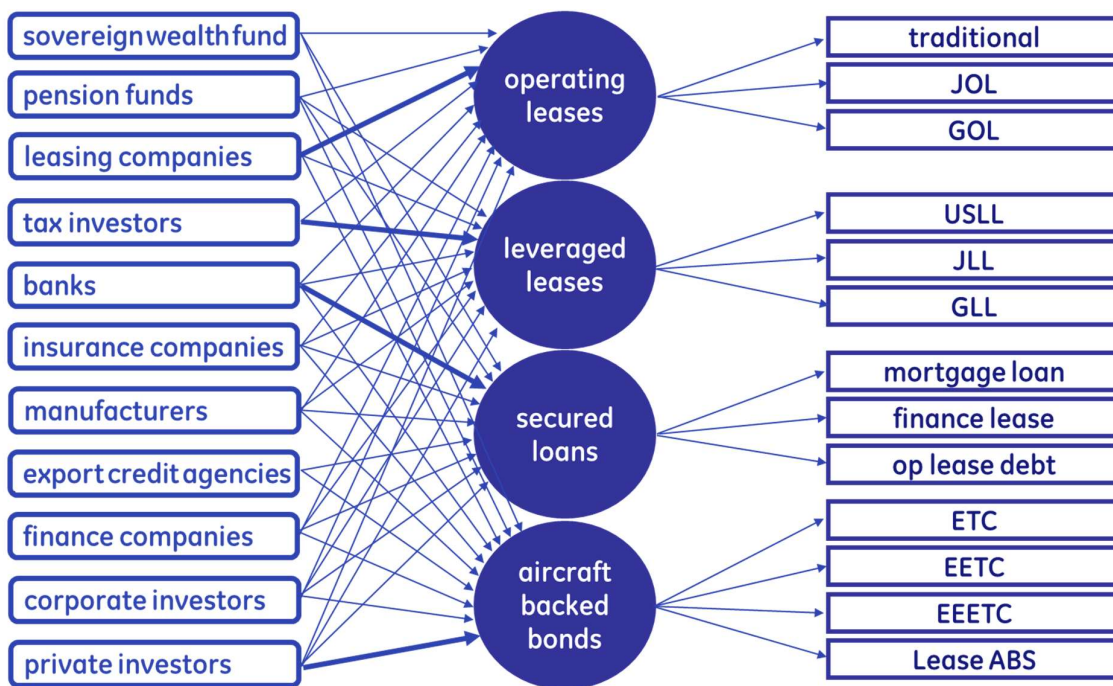
In the following, a couple of basic structures are discussed. There are endless possibilities to vary and combine, adding features, and structure the financings so as to obtain optimal risk-reward properties, accounting treatment, tax implications, flexibility, optionality, and contractual solutions.

Airline Loan
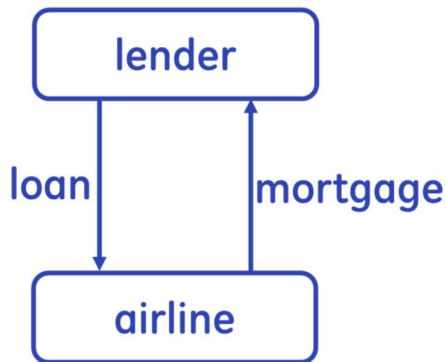
Operating Lease

Finance Lease
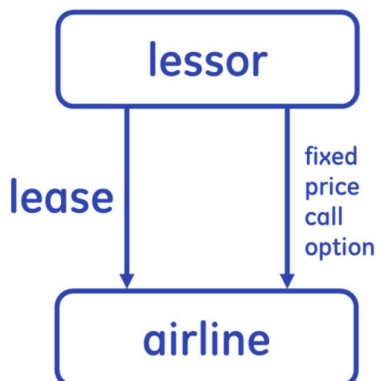
Airline Loan

The basic aircraft backed loan is directly to the airline, with full corporate recourse and secured by a mortgage over the aircraft as shown in the schematic to the left. Terms range from 3-15 years and the loans are typically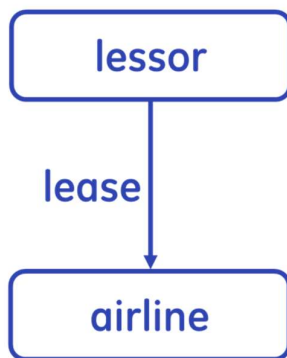 amortizing. LTVs range from 60% to 85%. The currency is often US $, lenders find comfort in lending in the currency that, by convention, is the currency in which prices of aircraft are quoted and traded. For carriers with little or no US $ revenue, this constitutes a currency mismatch. Baskets of currencies could well minimize the mismatch, both for the debtor and the creditor.

Operating Lease

The Operating Lease is a rental agreement where the airline (the Lessee) pays for the possession and usage of the aircraft over a contractual term. The periodic rent is meant to cover economic depreciation, interest and a return to the Lessor who is the owner of the aircraft at all times. Supplemental rent is often paid to cover maintenance expenses. The operator pays all other operating costs. The term is typically 3-12 years. At the end of the lease term, the lessee has to meet specific return conditions with regard to the maintenance status, so that the aircraft can be leased out again to the next operator.

Finance Lease

In a Finance Lease, the airline typically pays a sizeable first rental, the term is longer than for an operating lease, and the lessor has granted the lessee a fixed price purchase option at maturity.

In economic terms, the Finance Lease is equivalent to a loan with a non-recourse balloon, but differs in legal terms in that the creditor keeps legal title until the call option is exercised, rather than having a mortgage over the aircraft.

Non-Recourse Investor Loan
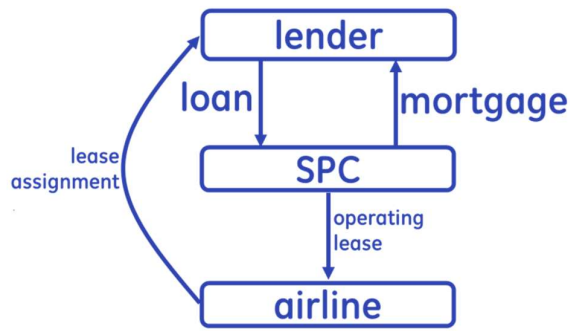
Pooled Investor Loan
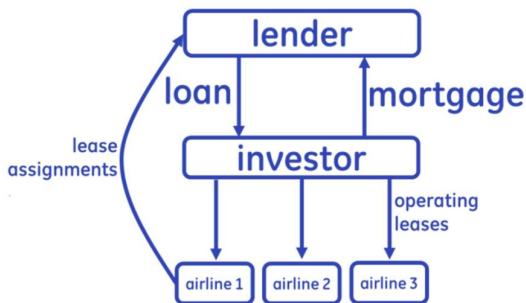
Syndicated Investor Loan
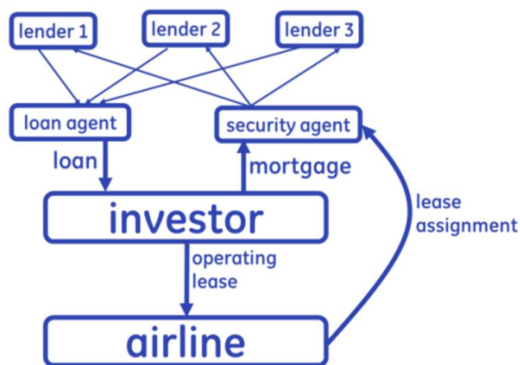
Non-recourse Investor Loan

An Investor Loan is a loan to the owner of the aircraft, but not the operator. Instead, the investor leases the aircraft to the airline. In addition to the mortgage over the aircraft, the lender gets an assignment of the lease as security. The cash typically flows directly from the lessee to the lender. When the borrower is a Special Purpose Company (SPC), the balloon payment effectively becomes limited recourse, and the lender looks to the value of the aircraft as the source of final repayment.

Pooled Investor Loan

In a pooled investor loan, the investor (or SPC) owns several aircraft on lease to one or more airlines. The advantage to the lender is that the diversification of aircraft and lessees reduces risk. The risk of all aircraft being run out at the same time and every obligor defaulting at the same time is remote. Cross-collateralization of the leased aircraft reduces risk even further.

Syndicated Investor Loan

The syndicated investor loan simply means that the loan is provided by several lenders. The roles of loan agent and security agent are sometimes split up in two.

In this case, the lenders rank pari-passu in the "waterfall", meaning that the proceeds from the collateral following an acceleration event or at maturity are split between the lenders according to their share of the loan , up to the amount of the claim. (See further the section on Default.)

Tranched Investor Loan
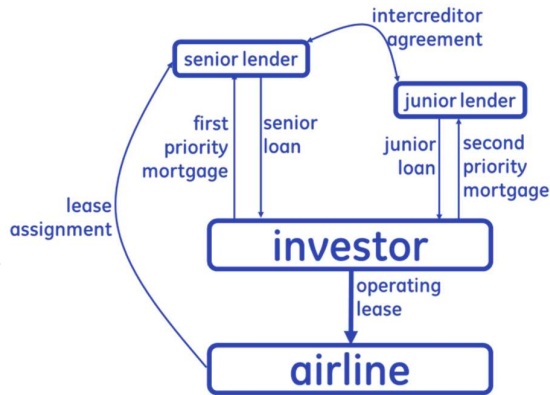
Tranched Investor Loan

The waterfall at Default could also be such that a senior lender gets all the proceeds until its exposure is fully repaid, whereas a junior, or sub-ordinated, lender only gets repaid when the senior is fully paid out. The junior position hence carries more risk, and warrants a higher return.

In any of these structures, the loan can take the form of a bond, which is easier to trade, and issued into the capital markets.

The repayment profile is sometimes referred to as the term structure. Aircraft backed loans are typically amortizing. The amortization profile can either be "straight-line", with equal amounts of principal paid back in each payment period, or "mortgage-style" or "annuity style" where the sum of principal repayment and interest is level. Obviously, any profile is conceivable, such as "high-low", 'low-high", "bullets" (zero principal repayment until maturity) or seasonally adjusted debt-service. When the final repayment at maturity is higher than the prior periodic payments, it is referred to as the "balloon". The balloon can either be recourse (it is the borrower's obligation to pay), or "non-recourse", or more aptly, limited recourse (to the asset). The non-recourse nature of the balloon is present when the borrower is a SPC, or in an airline loan, the lender grants the borrower a put option on the aircraft at a strike price at the balloon amount, or in a finance lease, when the airline has a call option on the aircraft. The pictures below illustrate the terminology.

Term Structure
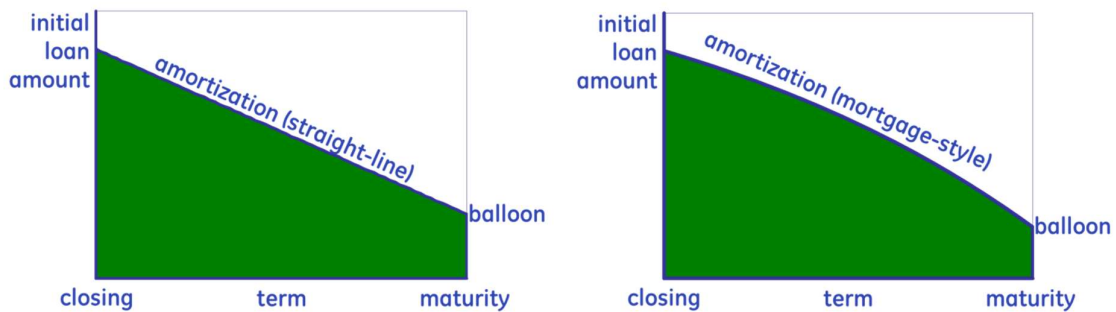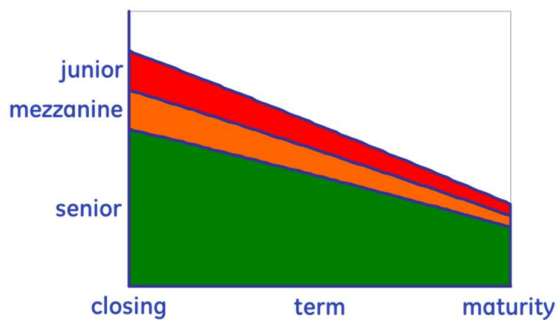
Default

What is Default?

When modeling the value or risk of an aircraft backed loan or lease, we need to estimate the probability of default by the airline over the tenor of the transaction.

A loan or lease contract will define **contractual default** in unambiguous terms, and specify the remedies available to the Lender/Lessor (Creditor). Events of Default typically involve late payments or the breach of various covenants (technical defaults). A contractual default gives the Creditor the right, after the cure period has lapsed, to accelerate a loan (thus becoming immediately due and payable) or terminate a lease.

A failure by the Borrower/Lessee (Obligor) to make payments as they fall due is a **payment default**.

After a default has occurred and not been cured, the creditor may exercise various remedies, including taking (or threatening to take) legal action and repossessing the aircraft.

It is sometimes difficult to establish if a late payment is due to INABILITY or UNWILLINGNESS to pay. There are a number of reasons why a late payment is not necessarily due to a general inability of an obligor to meet its financial obligations such as 1) inadequate administrative payment procedures, 2) commercial disputes between the creditor and the obligor, 3) a way to call the creditor to the negotiating table to restructure contractual terms, 4) an attempt to manage a temporary liquidity squeeze in the hope that the creditor will be patient or slow to take enforcement action.

If the creditor is vigilant in its collection effort, and the jurisdiction is well functioning, an obligor who is merely unwilling to pay will be forced to cure the payment default.

If the obligor is unable to meet its payment obligations in a timely fashion, it is illiquid. The obligor may then be forced by the creditors into bankruptcy and a potential liquidation of the

business. If the obligor has liabilities that exceed the value of its assets, it is technically insolvent. This condition may lead to illiquidity.

In many jurisdictions, an obligor that becomes insolvent may instead seek protection from its creditors in court for an interim period and reorganize the business within the allotted time frame. Some obligations may be met in full and others rejected. For example, under a US Chapter 11 Bankruptcy procedure, an airline may decide to affirm an operating aircraft lease, or to cure an aircraft backed loan. If not, the creditor will repossess the aircraft.

An obligor could attempt to reorganize outside court, in a negotiation with all or some of its creditors. For smaller companies with few creditors, this out-of-court procedure may be preferable as it would undoubtedly be faster and less costly than an in-court process.

A bankruptcy filing or an out-of-court restructuring of the obligor's liabilities must be labeled a **corporate default** in the sense that this can only be achieved if the court or the creditor group viewed the obligor as insolvent and truly UNABLE to meet its financial obligations in full. The corporate default is a state of the obligor rather than a specific loan or lease contract. This state is true even if some *specific loan defaults have been cured or lease contracts have been affirmed*. As the liabilities are restructured, the old obligations as they stood before restructuring are restated, modified or eliminated. Since everything is up for negotiation, under current market conditions, for modeling purposes, any continuing obligation must be seen as a new contract, even if the agreed terms remain unchanged.

When we try to estimate probabilities of default to be used in value and risk models, it is helpful to look at historical defaults. It should be straightforward to identify *ex-post* those obligors that have suffered a corporate default among large public companies. It is more difficult to do this for small and private companies. Historical payment records will help, but the distinction between payment defaults that were cured by solvent companies and corporate defaults by insolvent ones requires data that is sometimes unavailable or inaccessible.

In credit risk modeling, we have to model the probability of default *ex-ante*, and such an event must have a precise consequence within the analytical framework of the model. In almost all academic literature, the analytical consequence of a corporate default on a loan or lease is a stop in payments and a termination of the transaction. Corporate default is looked upon as an *absorbing state* – there is no way out.

When modeling asset backed loans or leases, we could assume that the consequence of default is a repossession of the asset, and in the case of a loan, the sales proceeds from the asset applied to reduce or discharge the loan balance, and in the case of a lease, the sales proceeds paid in full to the lessor. (A lender may force a sale of the aircraft. A lessor will either sell the aircraft following a repossession or, more typically, keep the aircraft and lease it to a new lessee.)

This is a simplifying assumption. In reality, as we have seen above, the reorganizing company may cure a loan default or affirm a lease and the transaction continues. However, it would be mathematically impossible in any multi-period model *ex-ante* to measure the post-default value or risk in a surviving loan or lease. There is no explanatory model that would help establish what the probability of default would be post-corporate default and post-reorganization. The question then is if this simplifying assumption of the model limits its validity.

In the case that the insolvent obligor reorganizes, and the particular loan default is cured or lease affirmed, the analytical consequence of this course of events is the same as in the crystallization of the asset position described above. Under a loan, it is reasonable to assume that the loan default will be cured if the aircraft market value exceeds the outstanding loan balance, but this outcome is equivalent to a repossession with a sale at market value that

discharges the loan in full – namely no loss. Under a lease, it is reasonable to assume that a lease will be affirmed if the contractual rent is lower than the prevailing market rent, but this outcome could be no more advantageous for the lessor than a repossession with a new lease on market terms – either way, no loss would occur.

Note that the specific transaction costs related to a repossession should not be counted in the case of cure or affirmation.

For modeling purposes, we assume that there are no arbitrage opportunities between leasing and selling an aircraft. If it was known that leasing would be a better alternative than selling, the market lease rent would go down and the market price for the aircraft would go up until the arbitrage opportunity had vanished. This does not imply that lessors cannot make money by leasing aircraft, it just means that in the world of a model, the lessor is able to cover costs and receive a fair compensation for the risk of owning the aircraft. The assumption of no arbitrage, in combination with an assumption that the lessee acts rationally (i.e. will only affirm the lease if the contractual rent is lower than market rent), means that the value of the aircraft with the re-affirmed lease attached is bounded upwards by the value of the aircraft itself, without a lease.

There are obviously examples where "under water" loan defaults have been cured and leases affirmed. Market information is imperfect. Asset markets are not always efficient – the operator may find that a specific aircraft has a higher value to them in use than the general market value (due to specifications, interior configuration or fleet commonality), and replacement assets may not be available without undue business interruption and transaction costs. But it would be difficult to analytically model this. A cure or affirmation of an "under water" asset may also be the result of a broader agreement between the creditor and the obligor involving other assets, or the extension of additional credit or providing additional assets. A single asset model can obviously not take those complexities into account.

Predicting Default

There are several ways to estimate the propensity of an obligor to default.

In an **expert system**, the estimate is a person's expertise and subjective judgment based on a set of key factors. This approach lacks consistency and objectivity.

In a **credit scoring** model, a set of key factors are identified that are believed to affect the probability of a default. Metrics based on the company's financial statements, management track record and processes, competitive position, market and regulatory environment and others are scored individually and weighted into an overall credit score or grade. By looking at scores of other obligors from the past and calibrating the weights with empirical default observations, using various regression techniques, we could estimate a default probability. Continuously updating the weights against new empirical data will make the system self-learning.

In a **structural** model, the risk of insolvency can be gauged by observing an obligor's equity price and volatility and thereby forecast the probability of the value of an obligor's assets (= equity + liabilities) becoming less than a given default barrier. The mathematics of this approach is related to pricing options, where the shareholders of a company can be seen as holding a call option on the company's assets. An example of this approach is Moody's KMV.

In **reduced-form** models, the default risk is observable in the market price of an obligor's traded debt instruments or credit derivatives. The return on the instrument can be split into a risk-free rate plus a risk premium. In the calculation of default probability, we would have to make assumptions about the recovery at default as well as the market price of risk. In a way,

the reduced form model is like running a debt-pricing model backwards and solving for default risk.

The two latter methods require that the obligor has securities outstanding that are publicly traded.

These methods can also be combined in various ways. Few airlines are public or rated by rating agencies, so lenders and lessors typically develop internal systems.

Default is relatively frequent among airlines. Figure 1 below shows ratings from Moody's KMV for public airlines between 1989 and 2009. The one-year Expected Default Frequency that KMV publishes has been mapped to rating buckets used by Standard & Poors. Around 75% of publicly listed airlines have been below investment grade (<BBB).
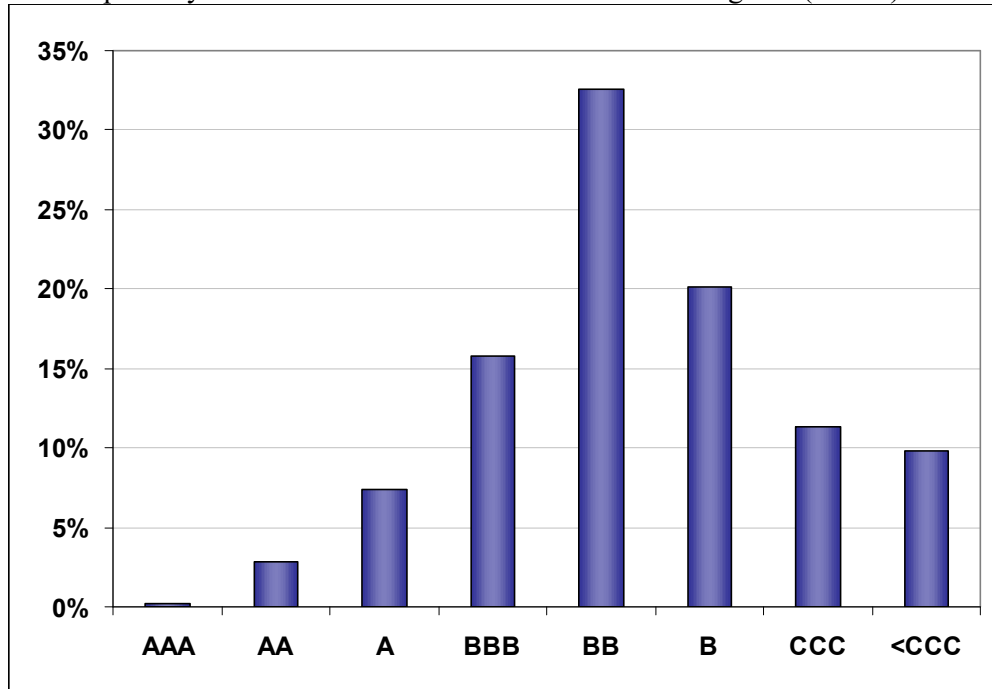
Figure 1. Histogram showing the distribution of S&P equivalent rating using monthly Moody's KMV EDF for airlines 1989-2009.

Some systems use the probability of default-within-one-year directly as a credit grade. Discrete grades can be mapped to one-year default probabilities. As aircraft backed loans and leases often have tenors of 5-15 years, the default probabilities must be estimated over the full term. There are several different ways to do this. In a structural or reduced-form model, the mortality term structure can be derived directly.

In credit scoring models, we can use empirical data. In Figure 2 below, an example of an average one-year transition matrix is shown. This matrix was published by a rating agency and covers data over 25 years. The center diagonal shows the probability that an obligor will keep its rating grade within one year, and the matrix also shows the one-year probability that an obligor will transition from one grade to another. As this data is collected across many different types of industries, there is no assurance that the matrix is representative for airlines.

FROM	AAA	AA	A	BBB	BB	B	CCC	default
AAA	91.6%	7.72%	0.48%	0.09%	0.06%	0.00%	0.00%	0.00%
AA	0.63%	90.4%	8.12%	0.61%	0.06%	0.11%	0.02%	0.01%
A	0.05%	2.15%	91.3%	5.78%	0.45%	0.17%	0.03%	0.04%
BBB	0.02%	0.22%	4.10%	89.6%	4.68%	0.82%	0.20%	0.31%
BB	0.04%	0.09%	0.36%	5.79%	83.3%	8.09%	1.03%	1.32%
B	0.00%	0.08%	0.23%	0.32%	5.88%	82.3%	4.77%	6.44%
CCC	0.09%	0.00%	0.35%	0.45%	1.50%	11.1%	53.5%	33.0%
default	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	100%

Figure 2. One year rating transition matrix compiled by a major rating agency for corporates. The matrix shows the empirical probability of transition from one rating grade to another.

From the matrix above, we have the one-year probability of default. If we assume that the intensity of default is constant over time, the term structure of the survival is an exponential decay. The probability of survival (=1-probability of default) is graphed below.
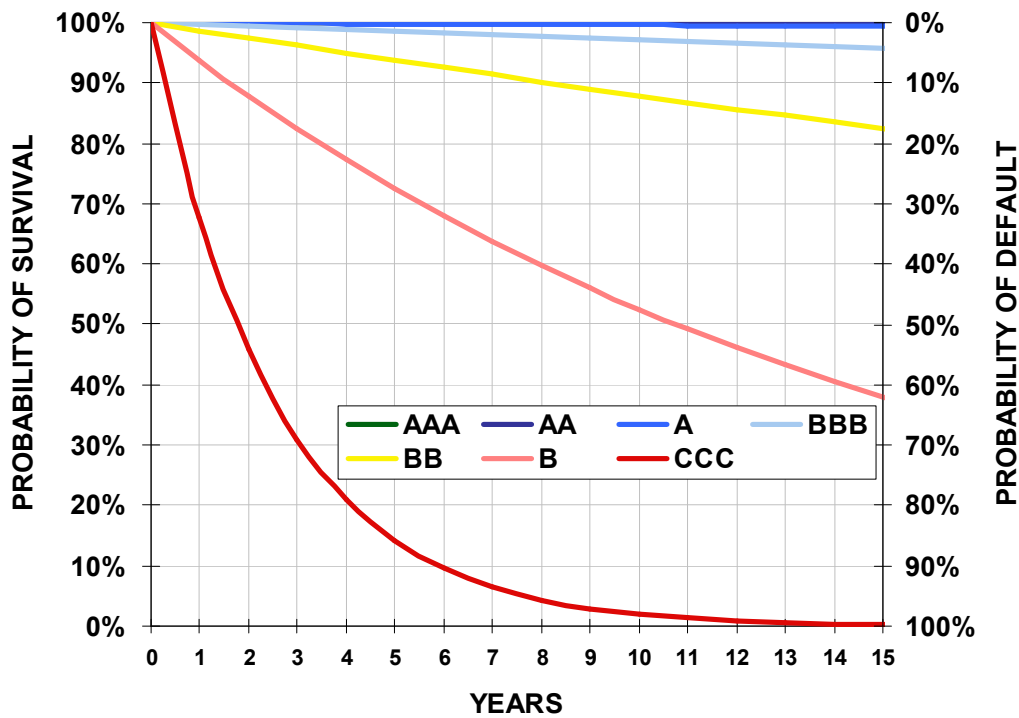
Figure 3. Survival curves for different rating grades assuming the one year default intensity is constant over time.

If we assume that the matrix remains unchanged year over year, we can multiply the matrix with it self to get to the two year matrix, and so on. We can then graph the cumulative probability of default over time (=1 − probability of survival), as shown in Figure 4 below.

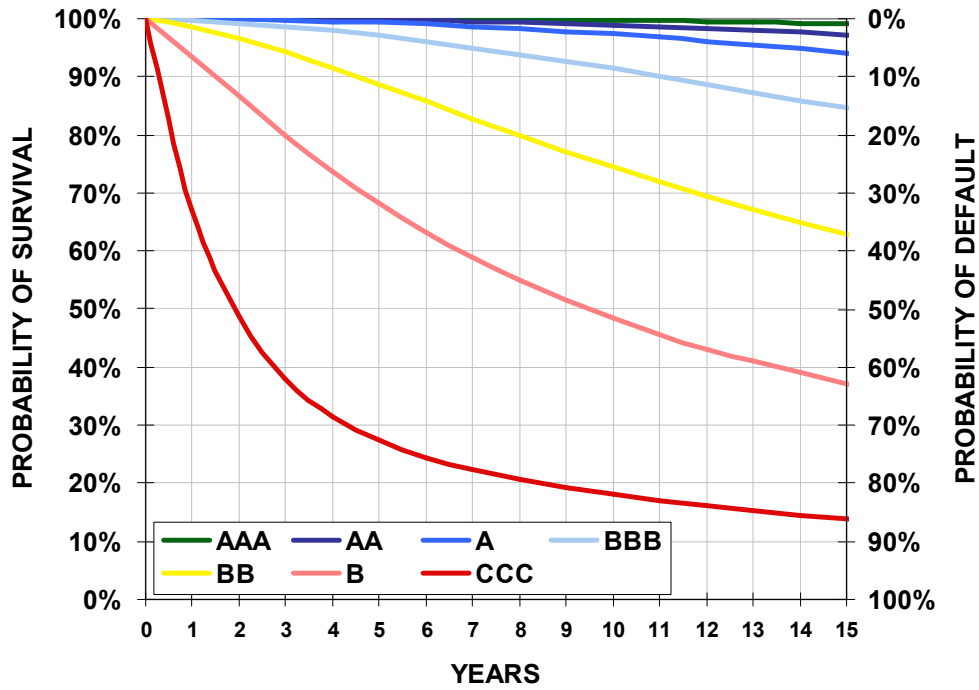*Figure 4. Survival curves for different rating grades assuming the one year transition matrix is constant over time.*

The same rating agency also published the cumulative average default rates directly, and the survival/default probabilities are shown below.
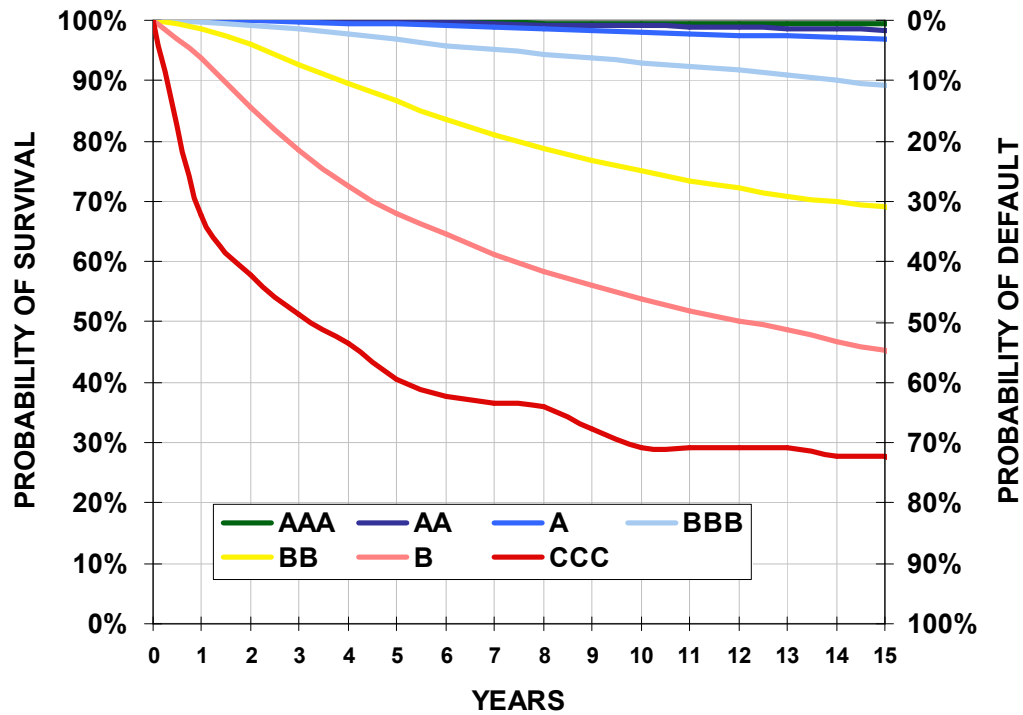

Figure 5. Empirical survival for corporates based on the same data set that was used for the rating transition matrix above.

19

The three methods give somewhat different results. The assumption of stable default intensity over time as well as the assumption of stable rating transitions are not really correct. We know that defaults tend to cluster around economic downturns. The empirical cumulative survival/default data is based on long term averages and may be outdated by the time enough data has been collected, and airlines may behave differently from other industry sectors. However, the charts give some idea of the shape of the survival term structure. We will address the issue of clustering in the section on the cycle.

How good are your predictions then? The challenge here is to compare a forecast that is a rating grade or a probability of default at various times to an outcome that is binary. Either you default, or you don't. One could track the average cumulative empirical default rate by rating cohort and horizon and verify how well this corresponds to the predicted survival/default term structure. But it would take many years and obligors to gather a meaningful amount of data over a long enough horizon.

A shortcut is to look at two separate aspects of the rating system:

1) how well does the rating system *discriminate* between the defaulters and survivors? and;
2) how well is the rating system *calibrated* to the overall level of default?

The power of discrimination over a specific time horizon can be measured by comparing the distribution of all survivors and defaulters by their assigned rating rank. Around a given rating level (the blue vertical line in Figure 6), we can measure the **probability of a true positive** (P(TP)) (the obligor survived and the rating was above the cut-off rating, the green striped area in Figure 6) and the **probability of a false positive** (P(FP)) (the obligor defaulted but the rating was above the cut-off, the red striped area in Figure 6).
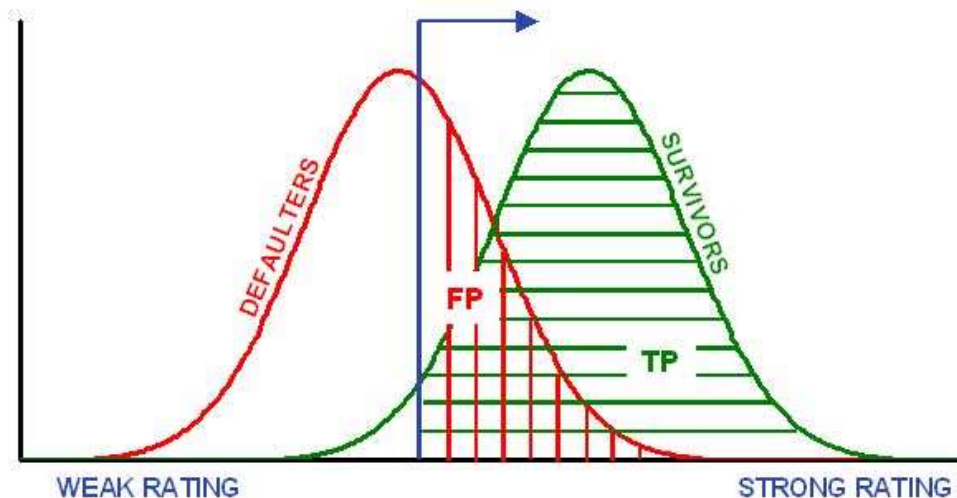
Figure 6. Distribution of rating grades one year prior to the determination point for defaulted obligors and non-defaulting obligors. Without overlap, the ratings would have discriminated perfectly.

We now move the cut-off rating rank along the rating axis, and trace a curve of the probability of a true positive against the false positive. We get the dark blue curve in Figure 7. The light blue line shows a random system without any discriminatory power. In a perfect system, the defaulter distribution would be completely to the left of the survivor distribution with no overlap and the curve in Figure 7 would follow the vertical axis and P(TP)=100%.
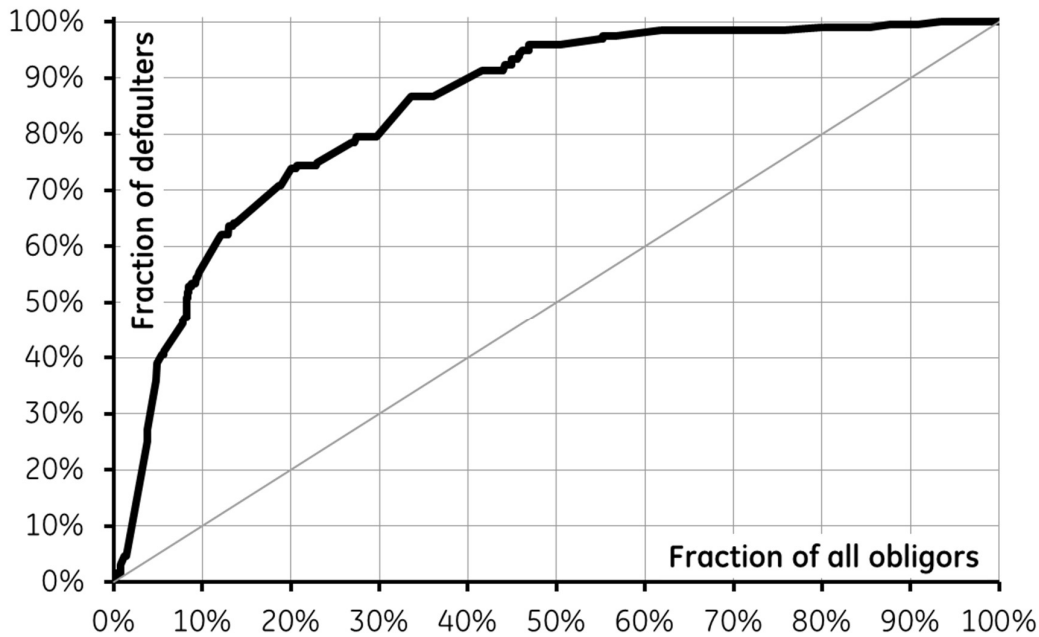
Figure 7. The ROC curve showing discriminatory power.

This curve is called the ROC curve (Receiver Operating Characteristic – from signal detection theory). The accuracy is expressed as the area under the dark blue curve, so 0.5 means no predictive power and anything above means better power, with 1 being perfection.

In a quantitative model, we must also assure that the projected probabilities of default are calibrated to actual levels. This can be done over time for different time periods; over a specific horizon; or on average, over a number of periods.
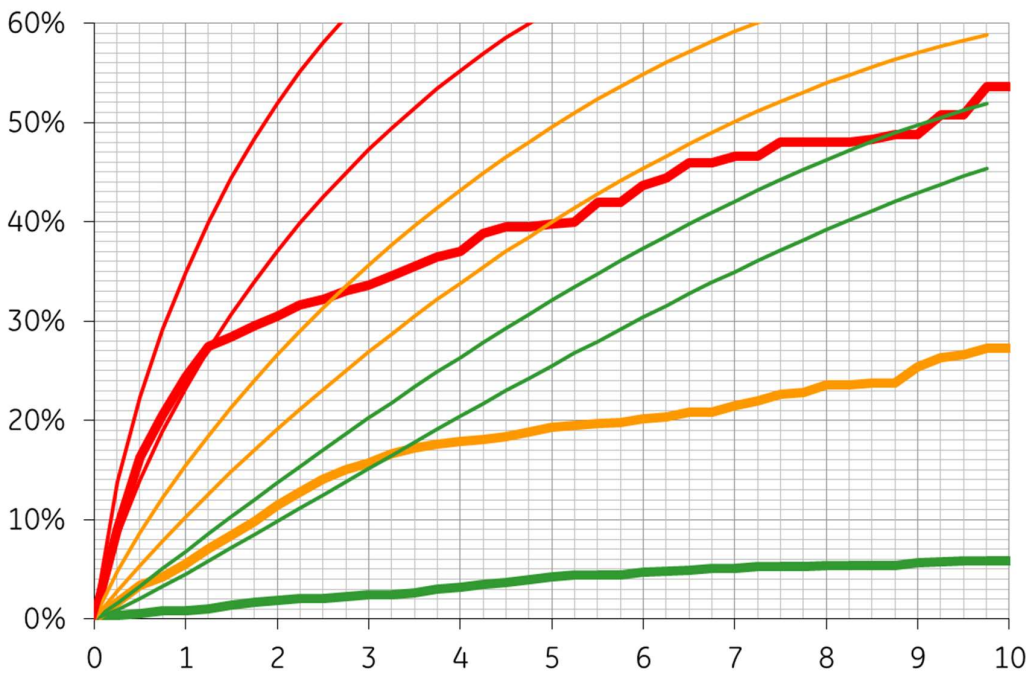
Figure 8. An example of miscalibration: The actual average cumulative default levels shown for specific rating buckets (thick lines) versus the predicted default levels for the same buckets, over the prediction horizon. In this example, the predicted default level was overstated, especially for the better rating grades.

Default is a binary outcome. If we assume that we will hold the loan or the lease to the end, it only matters if the obligor defaults or not. If we would try to model today what the value of the transaction could be at various points in the future, we would need to look at how the obligor rating could evolve over time.

The Cost of Default

A default event will entail costs for the creditor. The distress leading up to default will undoubtedly consume management attention. Additional time and resources will be spent when the default occurs, and most likely, additional legal fees and expenses will be incurred. In the event of enforcement, there will be costs associated with **repossession**, **repair** and **remarketing (rrr cost)**. Repair involves bringing the aircraft up to a technical condition to allow remarketing as well as obtaining or reconstituting technical records. This will often require maintenance to be performed earlier than would have been the case had the aircraft remained within the operator's fleet. Remarketing will involve the sale effort as well as bringing the aircraft up to the specification required by the buyer (or a subsequent lessee).

Apart from the resource costs, there is the capital cost for **downtime** during the delay between enforcement and final sale. This period varies greatly from jurisdiction to jurisdiction, and from case to case. The capital cost consists of interest carry and economic depreciation of the aircraft.

It is very difficult to predict what the actual costs will be in a default. The probability of non-cure/affirmation can be estimated based on jurisdiction and value of the obligor's contingent cure/affirmation option. The rrr cost is driven by aircraft type and jurisdiction whereas the downtime cost is often mainly jurisdiction specific. See section on Jurisdiction Risk on page 120.

Aircraft Values

Aircraft are depreciating assets and their values are volatile and uncertain. Both these characteristics are illustrated in Figure 8 below. It shows historical resale prices for commercial jet aircraft that are inflation adjusted and as a percentage of the estimated new price by age at the time of sale. The data was collected from 1974 to 1998. The solid red line is the trend line through the scatter.
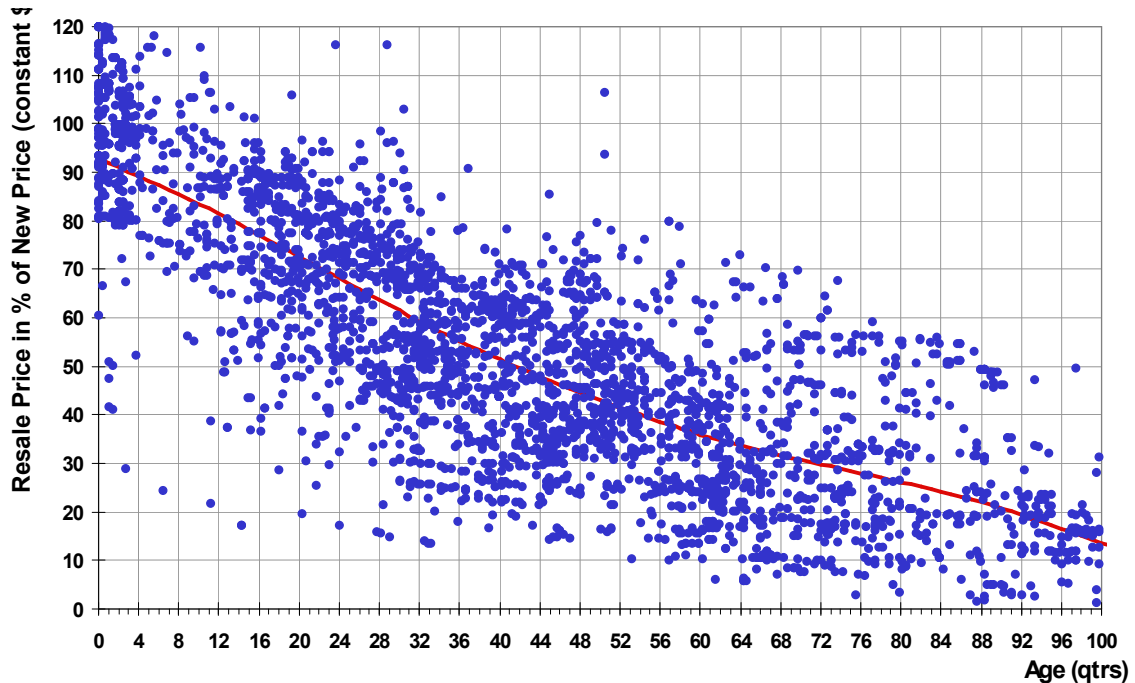
Figure 8. Scatter plot of constant dollar resale prices as a percentage of new price for jet aircraft collected 1974-1998. The dataset contains about 4,000 data points. Some sales may have had leases attached to the aircraft.

A similar picture is seen in the Figure 9 below. These are inflation adjusted and normalized appraised values from Ascend, an appraisal firm, for various jet types collected from 1974 to 2009. The thick blue line is the median with percentiles representing 1, 2 and 3 standard deviations.
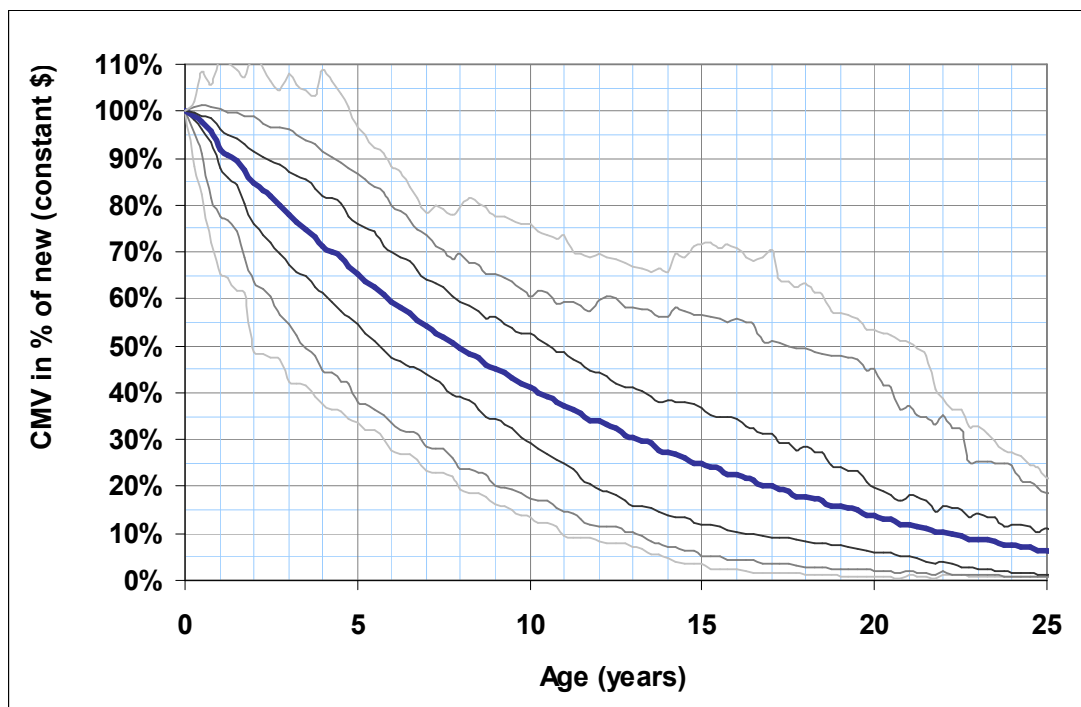
Figure 9: Constant dollar Ascend current market values as a percentage of delivery year value. The blue line is an average over all types and vintages. The gray lines shows percentiles one, two and three standard deviations out.

What Does Value Mean?

Before we go too deeply into a discussion of values, let us discuss what "value" means. For purposes of modeling an aircraft loan or lease, we need to understand what the *current market value* (CMV) is for the aircraft. This is a spot market value. A price that was contracted at a prior time is irrelevant for the value now. The aircraft is 'naked', in other words no lease, encumbrance or other contract is attached.

The CMV is different from a *base value,* which is some kind of "through-the-cycle" value, unaffected by imbalances in supply and demand or business cycles. Nobody trades at base values, but it may be a helpful concept to construct CMV projections. It is difficult to estimate base values in real time, it is really a long-term trend line that cuts through the ups and downs of the CMV.

CMV is not a distressed sale value. There are times when the sale-for-cash market is extremely illiquid, and actual sales prices are discounted heavily. In those times, a seller may need to offer vendor finance. In ISTAT's (International Society of Transport Aircraft Traders) definition of CMV, it is assumed that the asset is valued for its highest, best use. This is a somewhat dubious assumption since it is rare that all aircraft can always be expected to find a role in their highest and best use. In fact, older aircraft may well find gainful employment by flying very few hours – at peak periods or as back-ups. The higher operating costs may not matter much as long as the ownership cost is low.

For modeling purposes, we should think of CMV as the price a lessor could pay for an aircraft and expect to make a fair return on investment over the medium term. A lessor or lender always has the option to lease out rather than sell a repossessed aircraft if the sale-for-cash market is shut down. The leasing market is almost always more liquid than the sale market (albeit soft sometimes).

Bear in mind that a fair return reflects risk and the market price for risk fluctuates over time. In a normal and liquid market, this value should be close to what one could expect to sell an aircraft for provided that there is a willing buyer and a willing seller with neither under duress to conclude a transaction immediately.

Depreciation

An aircraft depreciates for at least two reasons. Firstly, an **individual aircraft** *deteriorates* physically over its life. The engines and airframe will require progressively more maintenance as they age, with associated increases in the cost of labor, parts and down time. An airframe will also pick up weight over its life and suffer from increased drag due to repairs and dirt. This causes escalating *absolute* operating costs. Secondly, the **aircraft type** will eventually face increased competition from more modern aircraft with superior economics thanks to improvements in performance, fuel burn, crewing, systems, aerodynamics, weight etc.

New aircraft types with superior performance may also come at lower prices (in real terms at least). As the manufacturers find that it becomes ever more expensive to build in another percentage point of improved operating efficiency, the attention turns to how to manufacture the same aircraft at less cost. Transitions from propellers to jets; analogue to digital cockpits; three to two people crews; lighter weight materials; low-bypass to high-bypass ratios; and metal to composites, are all *product* improvements making the aircraft more productive. Manual assembly to robots; riveting to bonding; and just-in time logistics; are all *process* improvements making the aircraft cheaper to manufacture. In both cases, the existing aircraft will lose out. This *obsolescence* will result in escalating *relative* operating costs. Both of the above will lead to a relative reduction of the cash flow generating capacity of the aircraft.

The theoretical value of the aircraft is the net present value of all the future cash flows that can be generated from the operation of the aircraft. Although aircraft do not have a limited *technical* life, their *economic* life ends when a positive net cash flow can no longer be generated.

If we take the PV of a cash flow stream that diminishes over time, we will get a similar shape to the charts above. In practice, it is impossible to forecast the cash flows from the operation of an aircraft type since different operators have different cost and revenue structures. We are trying to capture a market value for the aircraft, so the specific cash flow scenario of one operator is not relevant. But the theoretical analogy is useful to establish a basic shape of the depreciation pattern. It is also useful to measure for example sensitivity to fuel prices or fare levels in comparative analysis between different aircraft types.

An alternative measure of the aircraft's ability to generate cash flows would be to look at market lease rentals. However, the rental is only one of many parameters in an operating lease. The lessee's credit rating, tenor of the lease, the yield curve, jurisdiction, maintenance reserves, return conditions, early termination or extension options, lease deposits, and may other aspect will determine that value of a lease, so the idea of a *market rent* is somewhat illusory. But with a model of a lease, one could run the model "backwards" and solve for the current market value, given all the lease terms including contractual rent. That is certainly more meaningful than to inverse the Lease Rental Factor (the monthly rent divided by the market value of the aircraft – a very common metric in the lease industry with little real value) to find the current market value.

It is important to understand that between the two drivers of depreciation, one is related to the age of the aircraft and the other is related to the modernity of the type. Investing in aircraft delivered close to the end of their production runs has rarely been a successful strategy! An illustration of this is shown in Figure 10 below.
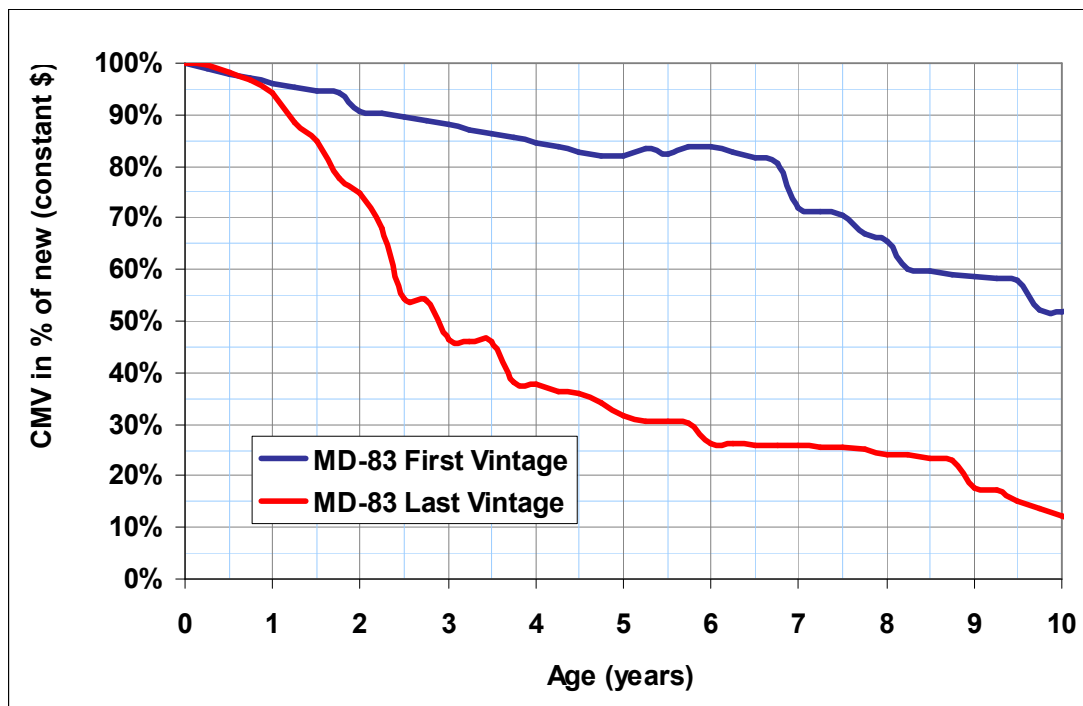
Figure 10. Relative constant dollar Ascend CMV for earliest and last vintage of the MD-83. Note that the early vintage has substantially better value retention.

Inflation

The charts above show value patterns in constant dollars. In reality, we face inflation. By inflation, we are not just talking about the escalation in aircraft prices, but about the erosion of the value of the currency. High inflation in the late seventies and eighties fueled the growth in aircraft leasing, when aircraft values sometimes kept rising in nominal terms. To forecast the value of an aircraft, we need to forecast the rate of inflation as well. (See section on Interest Rates).

A word on currency: by convention, aircraft prices are traded and priced in USD. That does not make an aircraft a "dollar asset" in the sense that hedging an aircraft on the asset side of the balance sheet does not imply a liability in USD.

There was a time when the USA had the lion's share of the world fleet, but today the USA represents less than 25% (in value terms) of the world fleet. If the USD were to devalue by, say, 50% against all other currencies over night, aircraft would become much cheaper to all buyers outside the USD zone, and demand would increase, hence, prices, expressed in USD, would rise. This would obviously drive inflation in USD. The impact of inflation on aircraft values is illustrated in Figure 11 below.
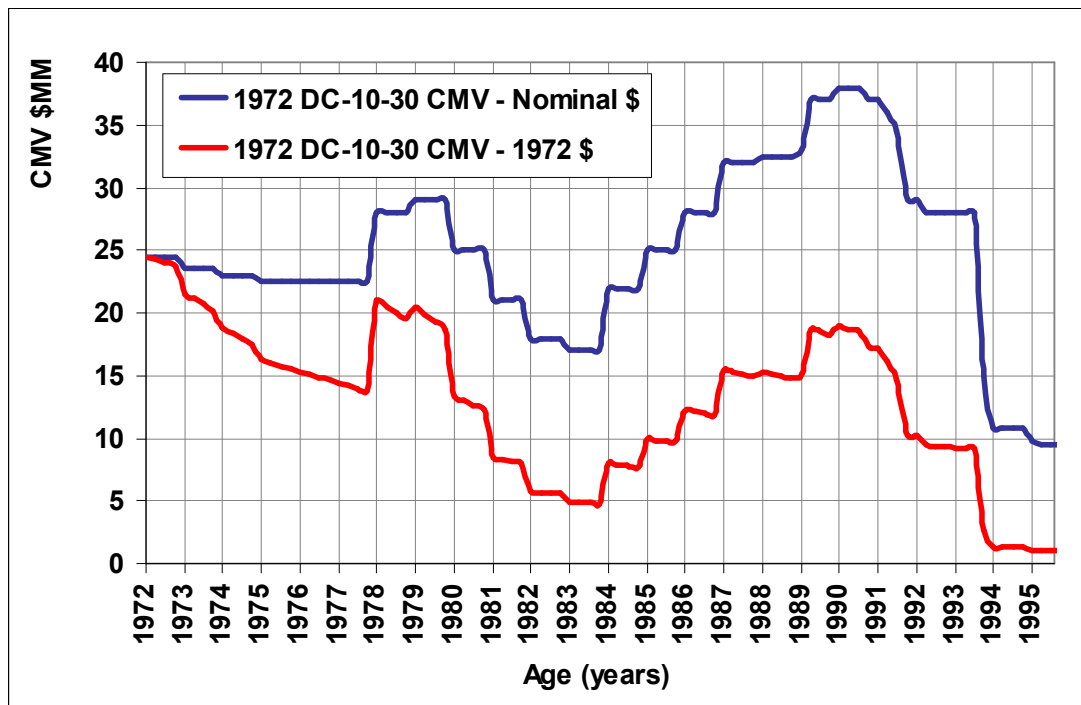
Figure 11. Ascend Current Market Value in nominal and inflation adjusted dollars. Note how big the impact of inflation was over this time period on value retention.

Maintenance Value

An important part of the value of an aircraft is its remaining maintenance status. As a new aircraft goes into operation, flight hours and flight cycles will be consumed against time- and cycle limited components of the engines and airframe. Major overhauls have to be performed within mandatory intervals. Subsequent overhaul, inspection, and restoration events tend to become ever more costly as the work scope and repair needs increase. The depreciation curve of a new aircraft will have a slope over the first three to five years reflecting the transition from a new aircraft to a "half-time" aircraft. The difference between a full value and the half-time value may represent over 15% for a new aircraft (somewhat theoretical as a new aircraft always has full potential, but still meaningful as a construction point for a value path projection), and increase significantly with age. For current aircraft types, we assume that maintenance value equals maintenance cost.

Airframe Maintenance

The airframe typically needs to be overhauled at set calendar time intervals. Some types are overhauled from zero time status to full-time status, whereas other types undergo "phased" checks, where some of the work scope is performed during intermediate events. Minor checks, sometimes called A- , B-, or C-checks can probably be ignored in the modeling exercise. The heavier check, the D-check, is significant. For newer generation aircraft, the checks are often referred to as time due, 8YE check / 12YE check. For simplicity, landing gears and auxiliary power units can be grouped with the airframe for maintenance value purposes.

The relative maintenance value can be illustrated in a "saw-tooth" diagram as per Figure 12 below.
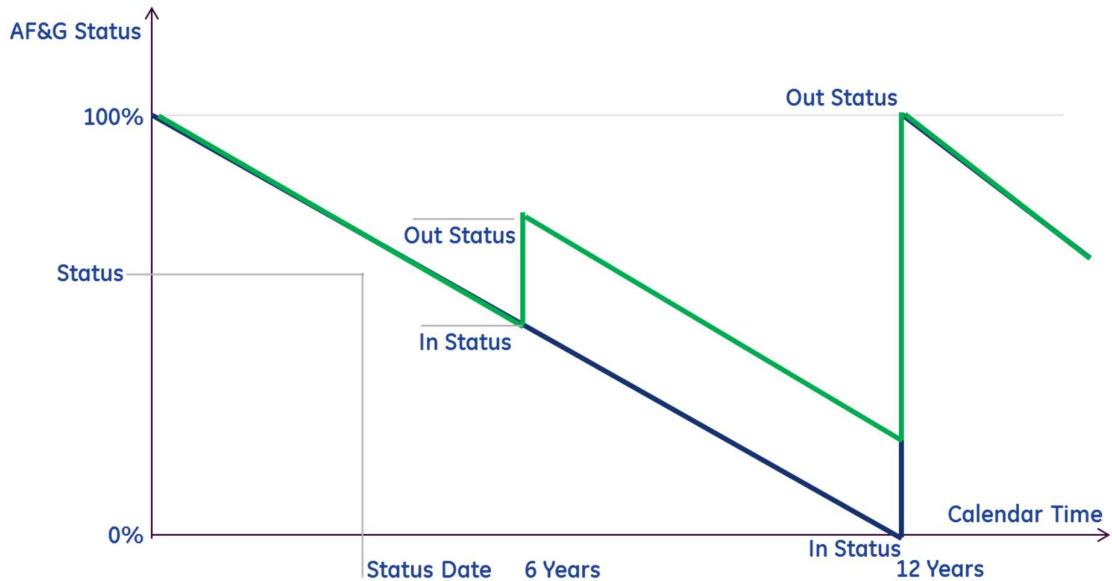
Figure 12: The maintenance value for the airframe and landing gear status illustrated for a type that has phased checks (green line), and one that has complete checks (blue line).

Engine Maintenance

Engine maintenance can be split in two parts: 1) Replacement of life-limited parts (LLPs), mainly discs and shafts that have hard limits, typically based on flight cycles, and 2) Refurbishment of the engine, repair and replacement of fan-, compressor-, and turbine-blades, guide vanes, casings, combustion chambers, bearings and gear box, to be repaired or replaced based on condition, and typically paced by flight hours and operating conditions.

Figure 13: A cut-out of the GEnx engine. Note the large fan and high by-pass ratio, and the discs with compressor- and turbine blades.

A modern jet engine often consists of four (or more) modules, the low pressure compressor, the high pressure compressor, the high pressure turbine and the low pressure turbine as shown in the schematic below in figure 14.
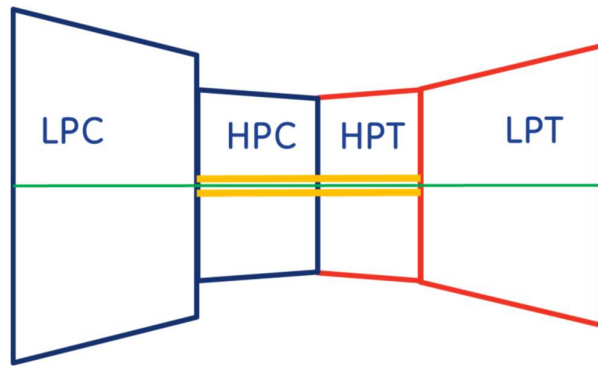
Figure 14: A schematic of the modules of a jet engine with the low pressure compressor (LPC), the high pressure compressor (HPC), the high pressure turbine (HPT), and the low pressure turbine (LPT).

The condition of the engine is monitored by measuring the temperature of the exhaust gas (EGT). An engine in good condition converts a large part of the thermal energy from the combustion of kerosene into mechanical energy, but as the blades get worn by impurities in the air, turbine blades oxidize, and cooling channels clog, the engine loses efficiency, and the exhaust gas temperature increases and approaches the Exhaust Gas Temperature limit. The time to reach this limit depends on operating conditions such as thrust settings, outside air temperature, take-off elevation, foreign object ingestion, and air quality. The operator expects a certain time (flight hours) on wing, but this is subject to some random variation. The removal of the engine from the wing, and the induction into a shop-visit can be predicated by either hitting hard cycle limits on the LLPs or approaching the "EGT margin", the primary measure of engine health. Depending on the engine make, some LLPs will need replacement, but others will have remaining life, and some modules will need refurbishment while others can be left to a future shop visit. It is an artful exercise to optimize the shop visits and decide what "build-standard" the engine should be overhauled to. LLPs with some remaining life could be replaced to avoid too short a time on wing until the next shop visit, and sold in the secondary market. The build standard should be optimized with respect to life time cost, financial considerations, and operational requirements. Typically, long haul aircraft hit refurbishment limits first, while LLPs is the limiter on short-haul aircraft.
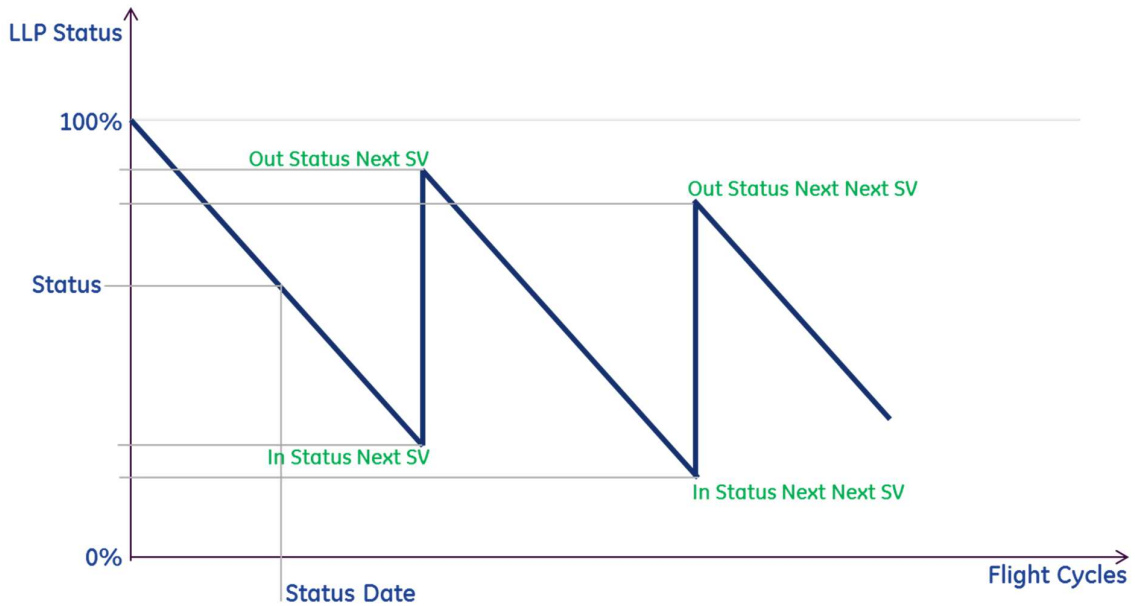
Figure 15: The maintenance value for the engine LLP status between various shop visits.

The flight mission in terms of flight hours per year, flight hours per flight, and the operating environment will determine the maintenance event planning. Obviously, this is subject to some variability, both among operators and for a given aircraft individual. Figure 16 below shows an example of the refurbishment and the LLP status of an engine.
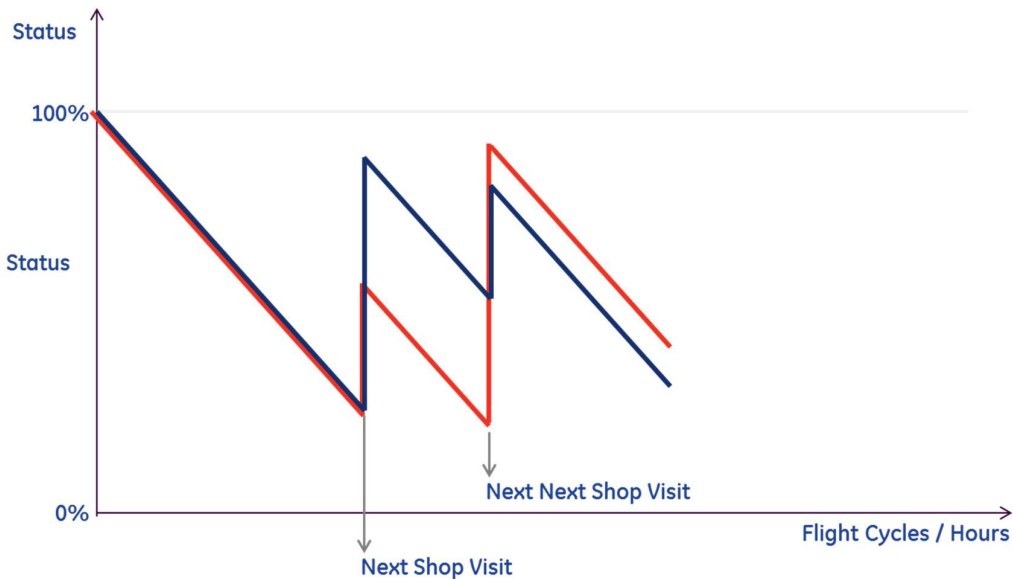
Figure 16: The first limiter will determine the time for the first shop visit of an engine. It is then a matter of optimization how to build up the engine for the second and third intervals.

For a young aircraft, it is reasonable to assume that maintenance value equals maintenance cost. However, as an aircraft ages, and an aircraft type matures, the market may become flooded with second hand spare parts and engines with remaining "green-time". A tell-tale sign is when the half-time current market value for the aircraft becomes less than 50% of the total maintenance cost for a full overhaul, and engine shop visit. At this point, it is clear that the maintenance *cost* no longer equates maintenance *value*.

30

Figure17. Scenario of combined maintenance value stemming from airframe overhaul, engine overhaul, and engine restoration.

Maintenance Reserves

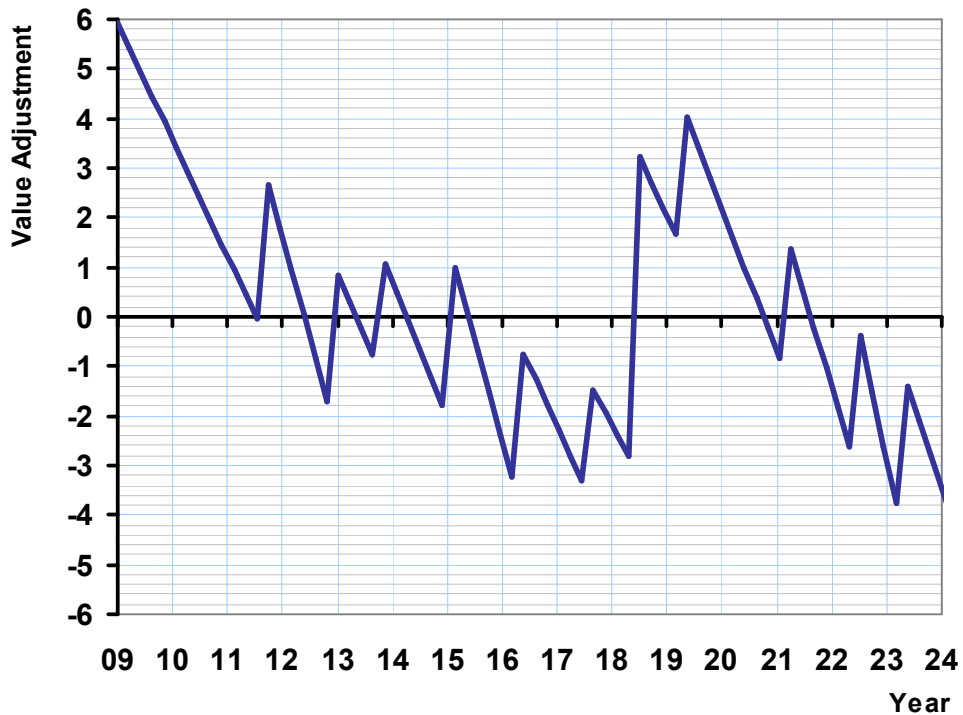Leases, and sometimes loans, often have provisions that require the operator to pay maintenance reserves based on months, flight cycles and flight hours to cover future maintenance costs. This is primarily a risk mitigant in case of default. To avoid issues in bankruptcy, such reserves are often substituted by "supplemental rent". These are normally collected on a monthly basis and on submission of the relevant claim, when certain defined work is accomplished, the operator will be refunded the cost. Some engine manufacturers or third party maintenance providers offer "power-by-the-hour" deals where the maintenance is provided against pre-agreed usage fees.

Lease agreements specify return conditions to assure that the aircraft has some remaining flying life before a next overhaul or shop visit. This also facilitates the transition to the next operator/lessee. In cases where maintenance reserves are not payable, the return conditions typically specify a restoration to delivery condition, or a monetary adjustment to compensate for the deviation from delivery condition. This is mainly available only to credit worthy airlines.

Scrap

Aircraft do not really have a limited technical life. Certification conditions may set absolute limits but these often get extended as a type ages and more operating experience is gained. Aircraft do, however, have a limited economic life. As the operating direct costs equal or exceed the cash generating capability of an aircraft, it will become economically obsolete. The last life in the aircraft is often the remaining flight hours and cycles under the maintenance program. After these have been expended, an aircraft will be scrapped. Components and parts may be used as spares, and the carcass goes to the smelter. In some cases, aircraft are scrapped

earlier, mainly because the parts for spares are in high demand. We have seen this recently for some short gauge variants of aircraft with more popular, larger variants in the family and with a high degree of parts commonality. We will revisit the 'death' scenario in the section on the cycle.

Volatility

The charts on resale prices and appraised relative value retention show, apart from depreciation, variation evidenced by the wide scatter of resale points and broad percentile bands. This variation reveals three principal features around aircraft values:

1) Some aircraft *retain their values* better than others;

2) Aircraft prices are not public and two transacting parties could settle for two different prices for two identical aircraft at the same time, hence we have instantaneous price *uncertainty*, and;

3) Aircraft values *swing* with changes in supply and demand over the business cycle.

This volatility is a key driver of risk in any aircraft lease or loan.

Value Retention

We can use history as a guide to see patterns of depreciation, but in the game of aircraft financing, the prize will go to those who can guess which aircraft types will best retain value into the future—and that is probably more an art than a science. Some characteristics that will be helpful to value retention of a type are *currency*: superior operating performance, lower operating cost, durability of structure, systems and components, mission capability well matched to market requirements, flexibility in configuration (cargo convertibility etc.), availability and cost of aftermarket support, regulatory compliance (safety, noise, emissions, and age related importation restrictions), and *liquidity*: number of units delivered, number of operators, order backlog, and dilution (absence of concentration). Aircraft that were delivered late in their production cycle will depreciate faster than those that were delivered early. The very early ones may suffer from excess weight and other early batch issues. However, when we forecast, although history teaches us that early deliveries should have better value retention, we will not yet know if a new type will display the currency and especially the liquidity as mentioned above. Therefore, new types may carry more risk than more mature types.

Price Uncertainty

Unlike stocks and bonds, aircraft are not traded on public exchanges with prices displayed in real time on computer screens. Aircraft trades are private and the prices are rarely disclosed publicly. If we ask 12 appraisers what the value is of an aircraft is today, we will get 12 different answers. They may have captured different data points, lack data points and make estimates, or infer from similar, but not identical, aircraft.

Almost all new aircraft and many used ones are not traded on the spot market, but on the "futures" market. The price is negotiated and agreed some time before the delivery, in many cases years ahead. Many aircraft are traded in bulk and would presumably benefit from a wholesale discount. Very often, the aircraft is sold with a lease attached, a financing commitment or guarantee support, with spare parts, warranties, credits, or other attachments that may add or detract value from the "naked" aircraft.

Aircraft may trade in a two-way exchange, and the price will not necessarily be reflective of the "naked" spot value. No two aircraft are identical, specifications or availability in time may be more or less suited for a specific buyer's requirement, all of which will affect the price a buyer is willing to pay. Appraisers typically state that their valuation assumes a willing buyer and a willing seller, under no duress. That assumption rarely holds. Again, this price uncertainty increases the risk in the loan or the lease.

Cyclical Swings

When air traffic growth is strong or when availability of aircraft is scarce, values of aircraft rise and vice versa. If the supply and demand would always be in balance, we could assume a kind of "base value" for an aircraft.

In reality, if we track historical aircraft values over time, it is evident that the values swing widely around the long-term trend-line. It is also evident that these swings are strongly correlated between different aircraft types although the amplitude of the swings will vary.

Larger and older aircraft tend to have larger relative swings. Larger aircraft (the size relative to other aircraft that are designed for similar range/capacity missions) are cycle sensitive because when the traffic demand is weak, they are harder to fill. When the traffic is strong, they are easy to fill and the operator can make lots of money. Figure below illustrates this difference with the example of a DC-10-30ER and a B737-800. The swings are measured as the deviation from a trend line that was calculated as a spline across the appraisal time series from Ascend. The B737-800 has swings of less than +/- 10% whereas the DC-10-30ER experienced around +/- 60% at the end of its economic life.
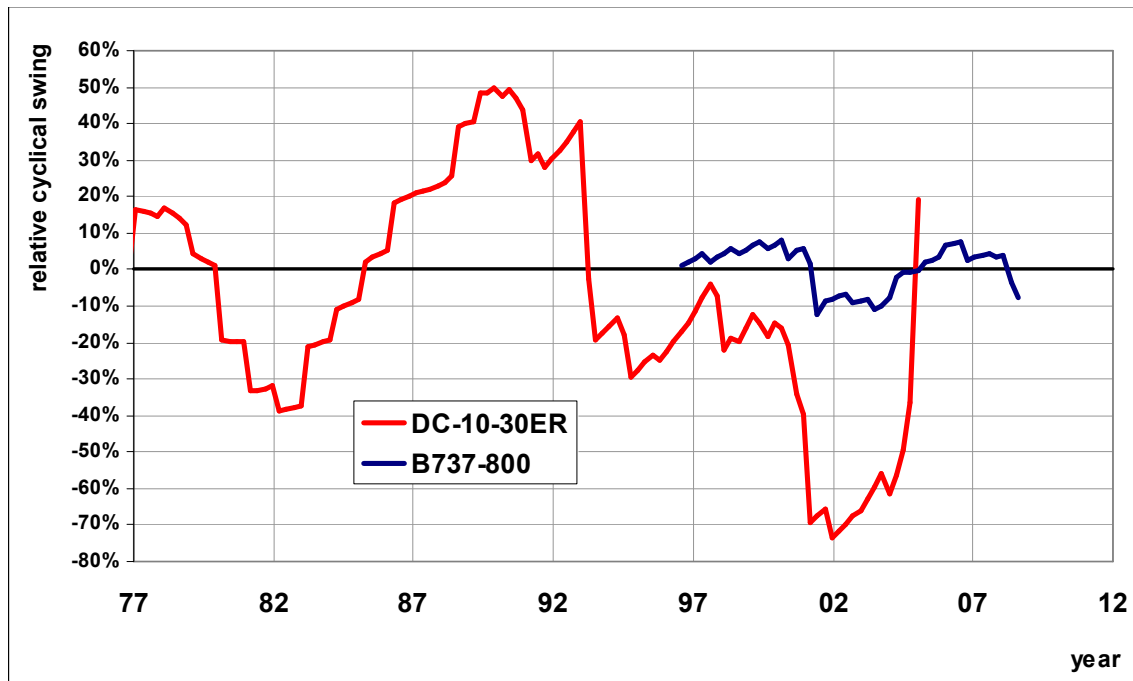

Figure 18: Cyclical swings around the trend line for two different aircraft types.

Older aircraft are cycle sensitive because when there is surplus capacity, the least efficient aircraft will not fly. When the capacity is scarce, older aircraft will still fly, especially on peak

time missions. The determinant is the ratio between capital cost ("have" cost) and operating cost ("use" cost).

Maintenance Status Uncertainty

We discussed the value of the remaining maintenance potential in an aircraft and that this must be part of the value projection. As the utilization of the aircraft, the exact interval times and the cost and value of maintenance work cannot be forecasted with precision, this uncertainty will be the source of further value volatility. Figure below illustrates this with the blue line showing the mean maintenance value and the thinner lines showing percentiles in a simulation.
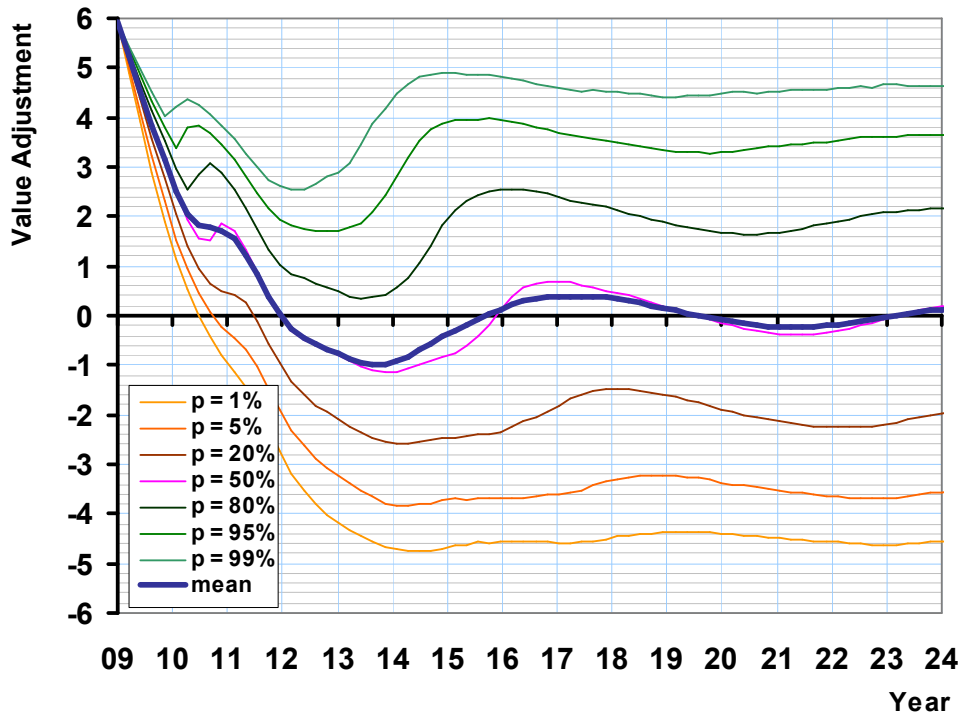

Figure 19: Simulation of maintenance value over half-life condition with percentiles reflecting uncertainty around timing and value of overhauls.

Forecasting the Value

Predicting the aircraft value out into the future is guesswork. History is a guide, but we have no assurance that history will repeat itself. Therefore, we need to break down the guesswork into components: base depreciation; inflation; maintenance value; and cyclical swings, each with the uncertainty that we have observed in historical data.

An example of base depreciation is shown in Figure below. Note that this theoretical construction line represents "half-time" maintenance value and that half the value of a full maintenance cycle should be added to the start value to get the initial value of the new aircraft. Note also how this straight line becomes a decay curve once the cash flow generating capacity of the aircraft starts to diminish. We can think of this point as the end of the "prime life" of an aircraft.

This is only one possible representation of the base depreciation curve. We would caution against using a *pure* decay curve. Although empirical data comes close to that shape when many aircraft types and vintages are blended, it fails to capture the difference between aircraft that are delivered early and late in the type's production cycle.
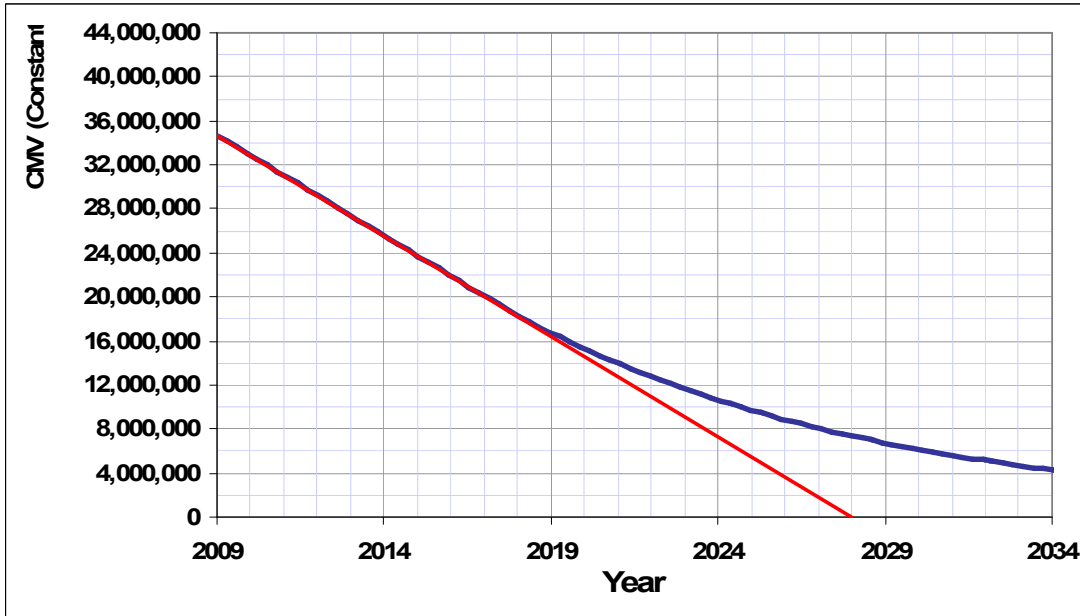

Figure20: Based on historical data, we could construct a base value path in constant dollars

We can simulate a path of the base value of the aircraft by combining (1) an expected depreciation function and (2) random movements to reflect the uncertainty. One way to generate the path is to introduce a normally distributed relative random movement of the value in addition to the expected depreciation or 'drift' in subsequent time steps. This is called a 'random walk' (discrete time steps), or Brownian Motion (continuous time). It may be appropriate to include a degree of "mean-reversion" in the random walk, also called a Ornstein-Uhlenbeck stochastic process. An example of a path is shown in Figure .
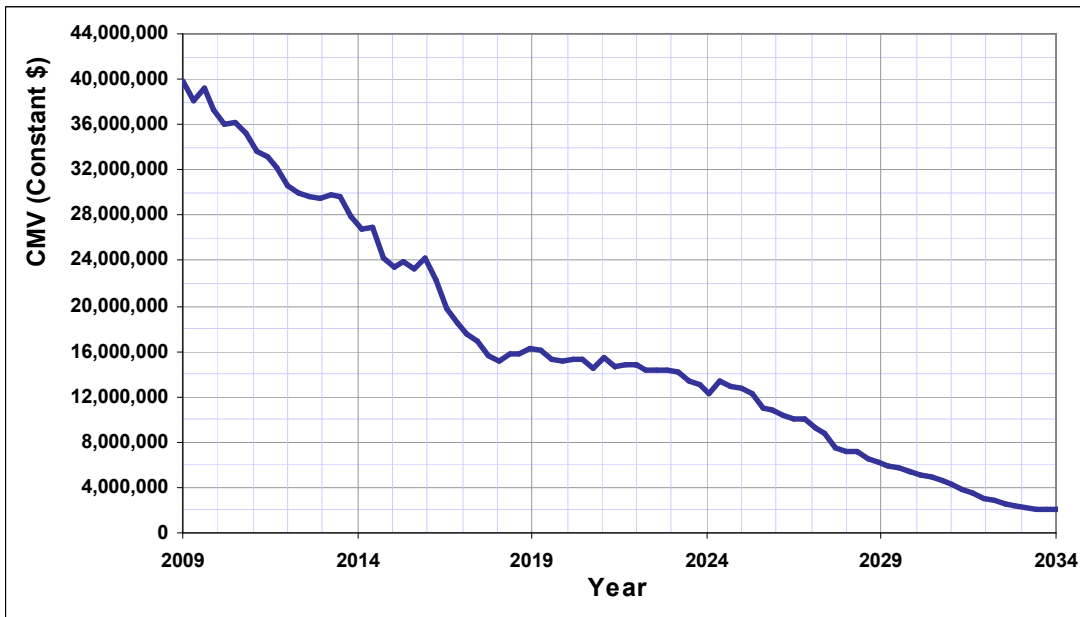
Figure 21: shows one possible path for the base aircraft value using the expected base depreciation above and a random walk.

We can generate a multitude of possible paths with the help of a computer program.
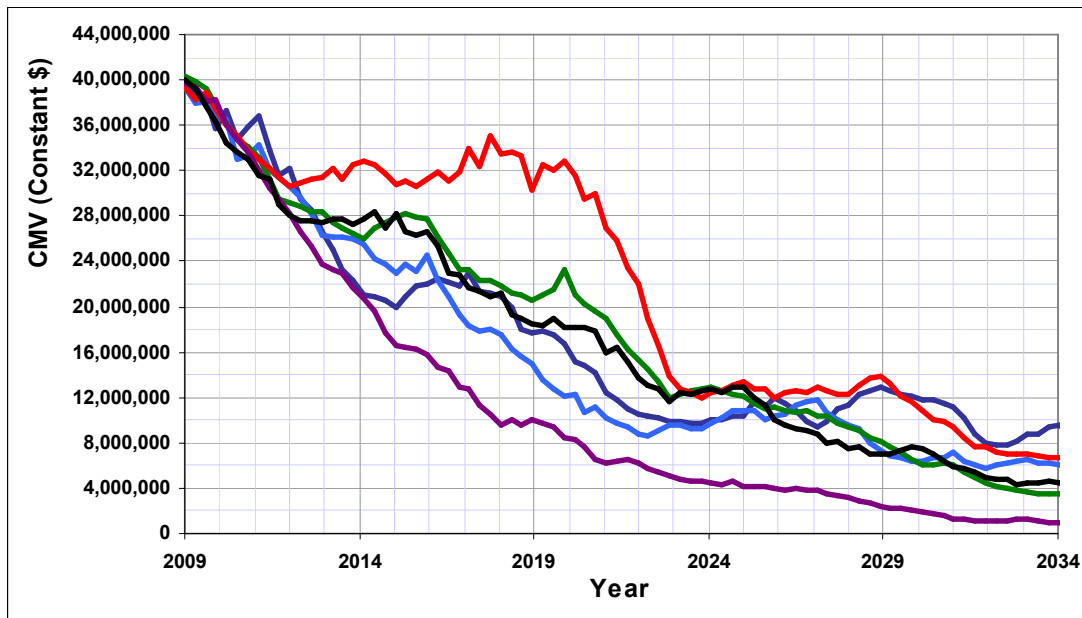
Figure 22: shows six random paths.

This Monte Carlo simulation can generate a desired number of possible paths that will reflect both the desired expected value and the desired probability distribution around this average.

Figure 23: shows the average of 500,000 random paths (dark blue line) and the median with percentiles representing 1, 2, and 3 standard deviations (in a Normal distribution), respectively.

Note that the picture is similar to the historical average and dispersion that we have observed.

We can now add a projected level of inflation. We will discuss inflation briefly under the Interest Rate section.
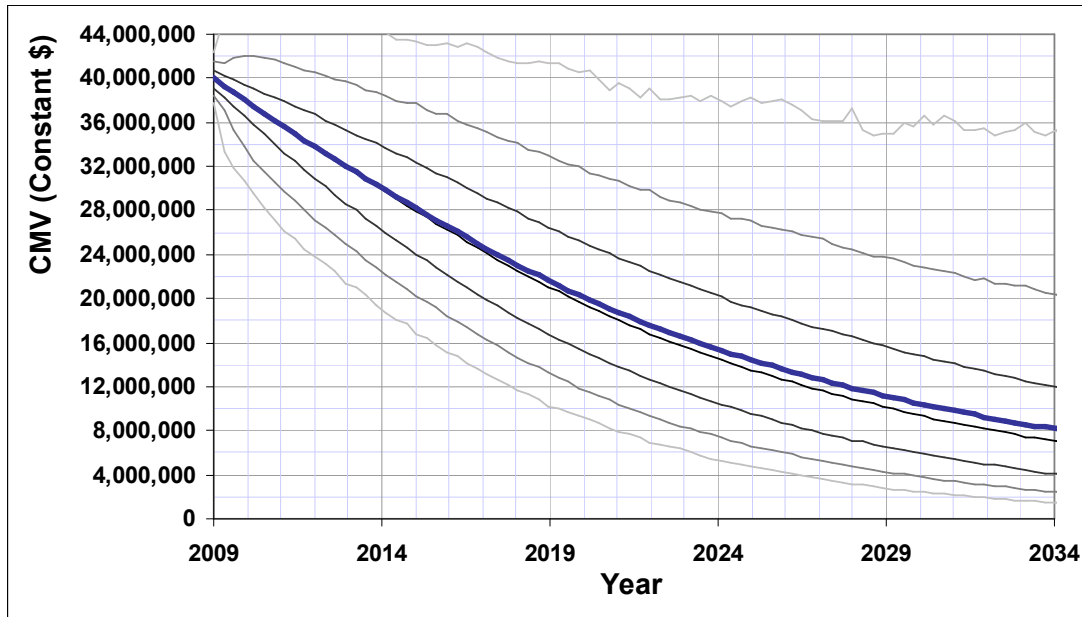
Figure24: Same as Figure 23 but with inflation added.
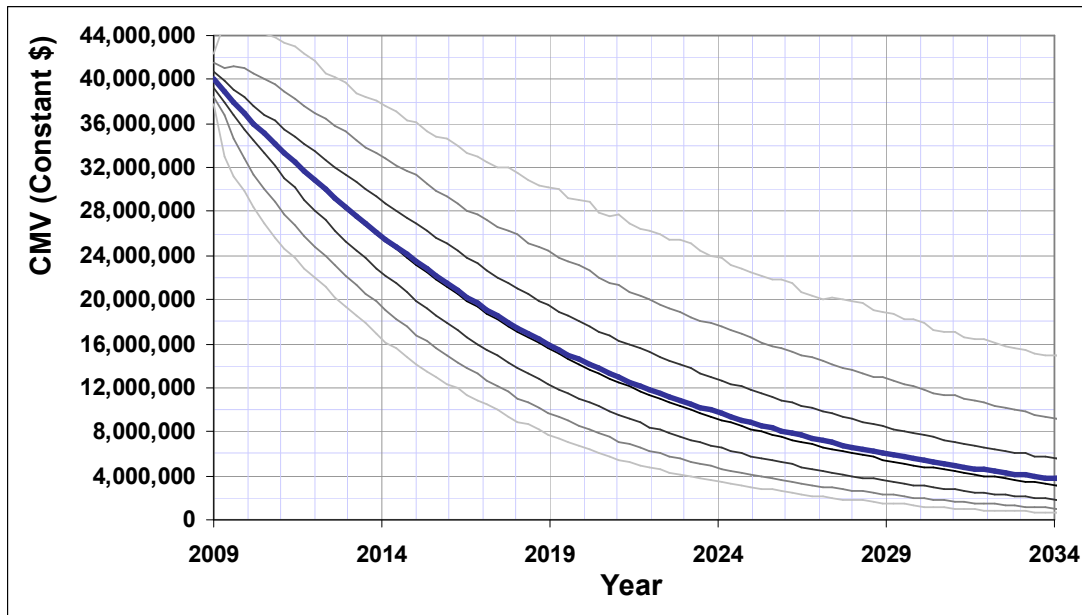
Forecasting Accuracy

Although historical depreciation and volatility of aircraft values is useful when forecasting, the acid test is how well the forecasts stack up against values that we measure ex-post for different forecasting horizons. Figure below shows the ex-ante projected future value for a specific aircraft against the appraised ex-post current market value.

Figure 25: Ex-ante forecasted value over a 6 years horizon against ex-post appraised value.

By taking the difference between the projected value and the appraised value, and dividing by the projected value, we get the **relative forecasting error** for different forecasting horizons. Figure 26 below shows the percentiles for 50%, 84.1% and 15.9% quarterly over a 7 year forecasting horizon for a sample of projections made by PK AirFinance during the period 1995 to 2009. The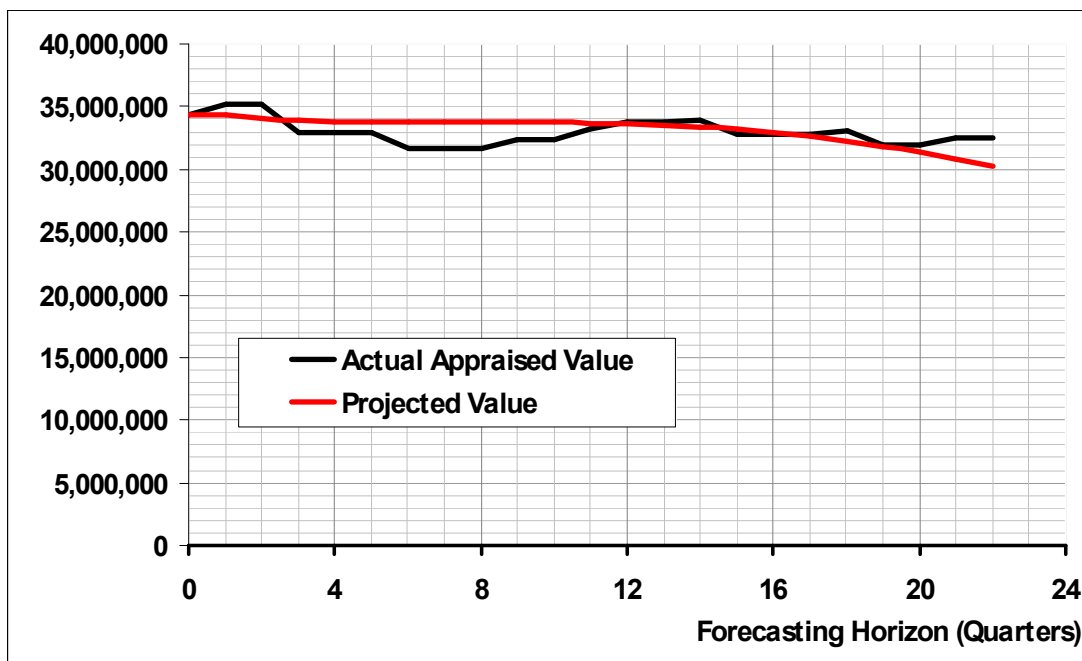re is a total of 886 projections of which 113 have spanned the full 7 years. Median values were overstated by 15% at the 7 year point.

The dispersion of the forecasting error is illustrated by the 15.9% and 84.1% percentiles (these percentiles correspond to ± one standard deviation in a Normal distribution).
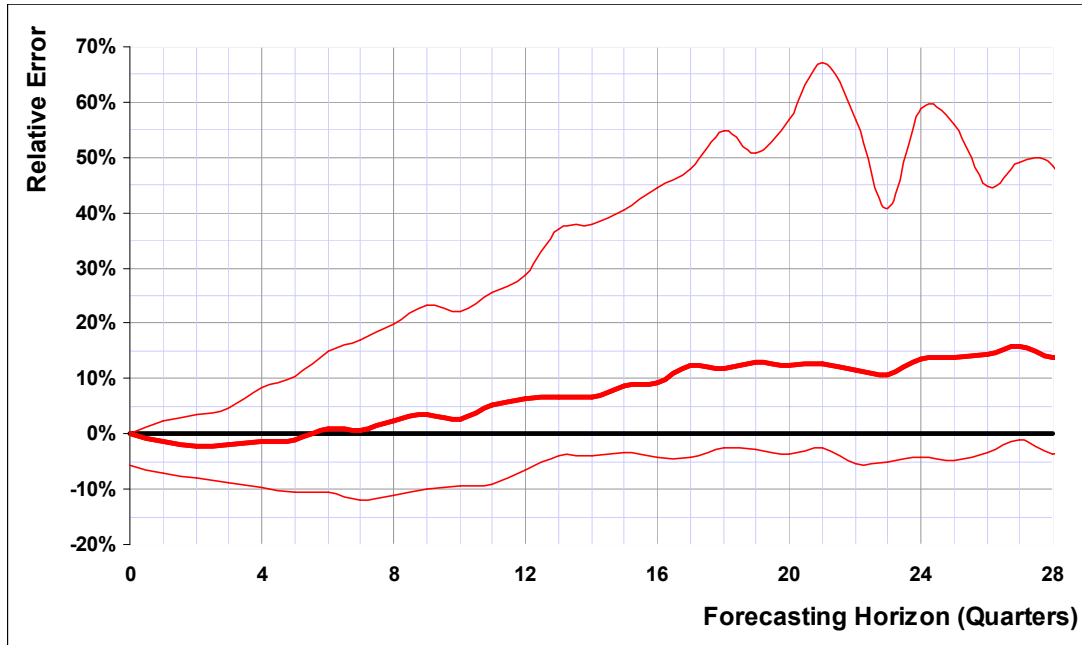

Figure26: Relative aircraft forecasting error observed on a sample between 1995 and 2009.

A considerably better record is obtained if only in-production aircraft types are included in the sample, as can be seen in Figure.
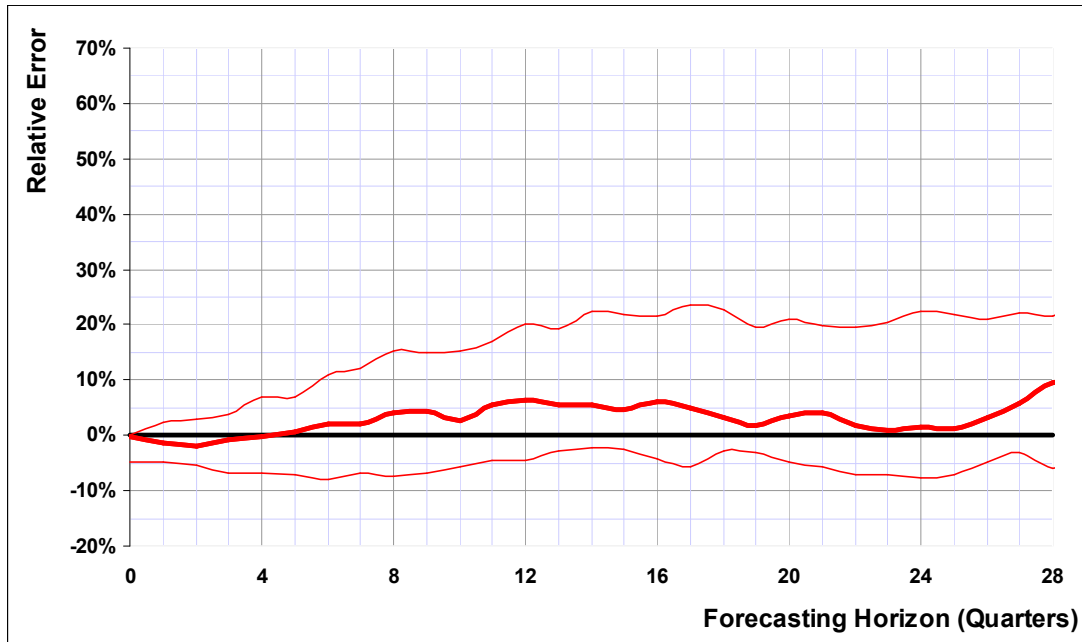

Figure 27: The same sample as above but with only in-production aircraft types.

38

Switching Costs

When an aircraft is repossessed or returned under a lease, and we look for a new operator, buyer or lessee, the aircraft will have to be configured to suit the new operator. The cost for this can range from zero to tens of millions of dollars. If the aircraft is bought to be scrapped, no configuration modification is necessary. The existing interior, in-flight entertainment system, galleys, etc. may be perfectly suitable for the next operator. Or, the new operator may need the aircraft to be configured to exactly the same specification as all other aircraft of the same type in its current fleet. For narrowbody aircraft, this cost may be up to $2.5mm, for widebodies, the cost may be $10mm or more. It is sometimes unclear what the appraiser means by 'current market value" – is this what one could expect to sell the aircraft for before or after having incurred the cost of a reconfiguration? One could be pragmatic when modeling switching costs and assume that it is a cost the creditor will incur, specific for each aircraft type, but that that this cost is capped at a set percentage of the aircraft value in each scenario. If we take an Airbus A330-300 as an example and estimate that the representative reconfiguration cost is $10mm, but that the cost is capped to 30% of CMV, in a scenario where the aircraft is worth only $20mm, the reconfiguration cost would be limited to $6mm. In other words, we assume that this aircraft would be older or that we face a depressed market, and that a buyer would settle for an inexpensive configuration. Such a percentage cap would not impact newer aircraft or aircraft traded in a buoyant market.

Other switching costs would be bridging checks, search costs, ferry, transaction costs, downtime, etc.

Interest Rates

We are interested in interest rates for two reasons:

1) fluctuating interest rates over the term of our loan or lease will affect the **value of the transaction**, and

2) interest rates are correlated to **inflation**, something that will directly affect the future nominal value of the aircraft.

What is Interest?

Interest is a fee paid for borrowed money. It covers 1) **time-value of money**, the amount a lender requires to be indifferent between having $100 today versus $100 at a later date, say in a year, and 2) **credit risk**, the compensation a lender requires for the risk of not being paid back. In the following, we will look only at the time-value of money component, where the Interest Rate is the **risk-free rate**.

Some Terminology

The **periodic rate** is the interest that is charged for each period, divided by the amount of principal. The **nominal annual rate** is defined as the periodic rate multiplied by the number of

compounding periods per year. The **effective annual rate** is the equivalent rate if annual compounding were applied. The nominal annual rate together with the compounding frequency will determine the exact interest stream. However, to allow comparison between different loans, the effective annual rate is preferable.

The **calculation basis** in the bond markets is a month deemed to have 30 days and a year of 360 days, whereas the Libor market (London Interbank Offered Rate) calculates actual days elapsed between payments, and the year is deemed to have 360 days. The **duration** is the time elapsed until the principal amount is due to be repaid.

In a **bond**, the issuer pays the holder principal and a coupon. In a **loan**, the borrower pays the lender principal and interest. The two are absolutely equivalent, except that a bond is typically a tradable security. The interest rate can either be **fixed** for the full term or for each interest period as a **floating** rate.

In the following, we will discuss fixed interest rate instruments. A risk-free bond would be assumed to pay the risk-free rate. There is a debate as to what the risk-free rate is. For practical purposes, we will assume that the risk-free rate is Libor for floating rate and the swap rate for fixed rate. If the rate is fixed, the whole cash flow stream of the loan is fixed.

The Yield Curve

The interest rate depends on duration. The relationship between the market risk-free interest rate and duration can be expressed as a yield-curve. Figure below shows the yield-curve for two different dates. Note that the yield curve typically has an upward slope, but sometimes the slope is inverted.

Note that both the compounding basis and the calculation basis need to be defined and consistent across the durations.
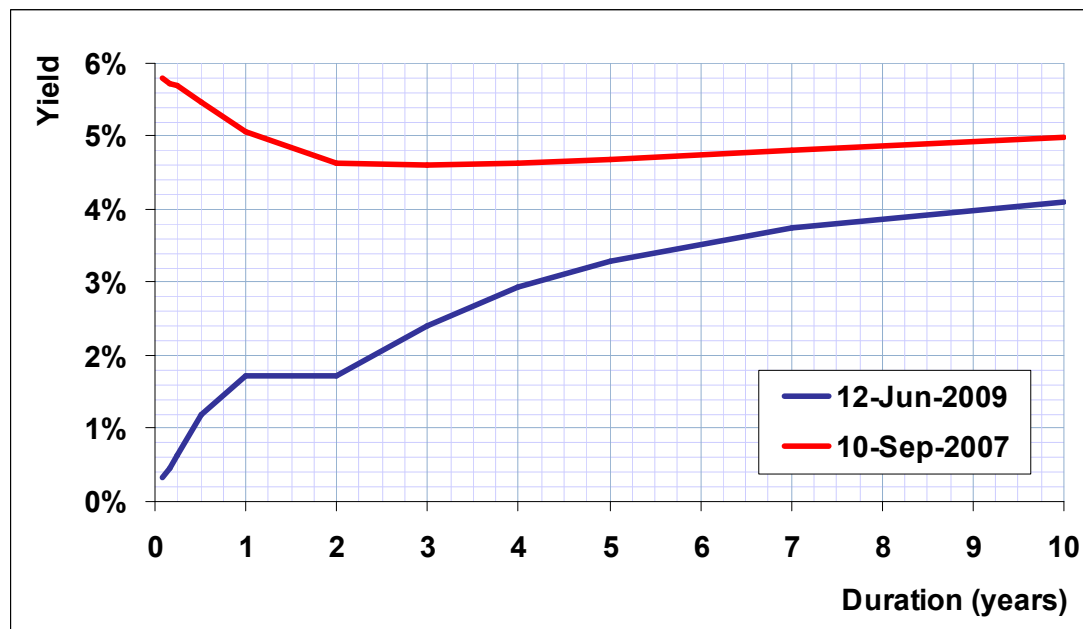

Figure 28: The yield curve for Libor and swap rates on two different dates.

The Value of a Risk-Free Bond

The value of a risk free bond is the present value of the cash flows from the interest coupons and the principal redemptions, discounted at the risk-free rate for the duration of each payment.

At issuance, assuming the bond yield is identical to the risk-free rate, the value of the bond will be its nominal value. Over time, the yield-curve will move with the market, and the risk-free rate for the remaining duration will now be different from the nominal yield of the bond. Hence, the present value will differ from the nominal amount of the bond as the cash flows will be discounted using the new market yield curve. We can refer to this as the mark-to-market value.

A bond will trade at a discount if interest rates are *above* the issuance rates, and at a premium if the interest rates are *below*. We are simplifying here in that we are not allowing for liquidity effects.

Amortizing bonds

In aircraft financing, the loan or book value is typically amortizing. In order to determine the risk-free interest of an amortizing bond, we have to look at the amortization profile. The interest cash flows have to be calculated for each over-layered strip, based on the duration. Figure below shows an example of an amortization profile.
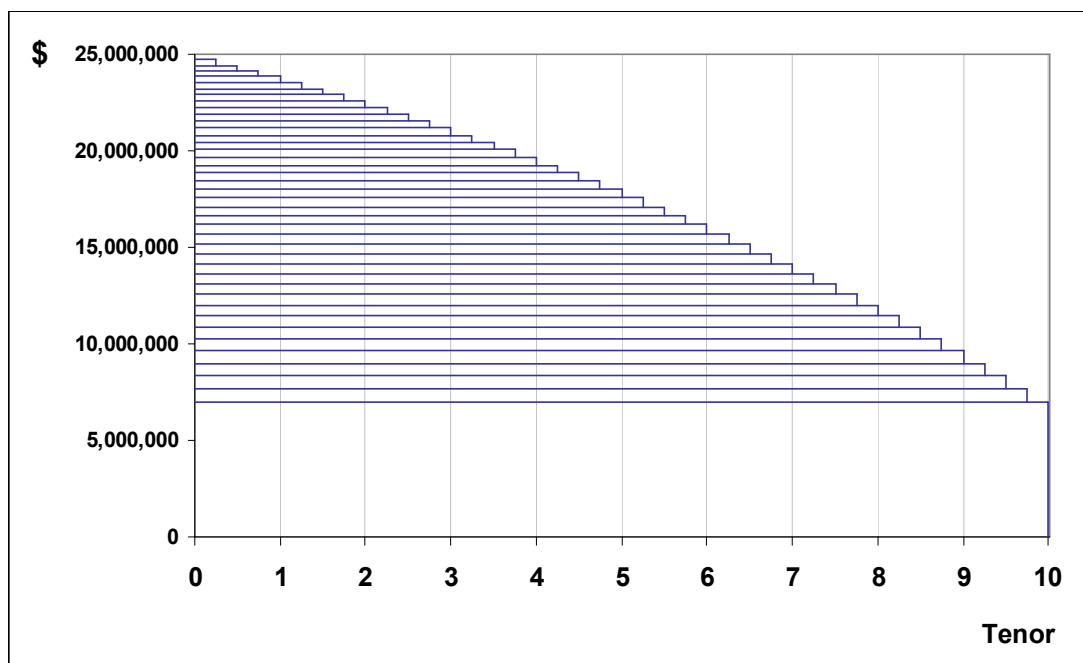
Figure29: Outstanding principal of an amortizing bond.

In this case, the amortization is called "**mortgage-style**" where the sum of principal and interest (debt service) is constant for every payment date until the balloon payment at the end of the term. Another variant is the **straight-line amortization** or **level principal payments** where the principal component is constant.

To find the rate for durations between two points on the yield-curve, we have to interpolate. To set a fixed contractual interest rate for the bond, we calculate the rate that will yield a series of cash flows with exactly the same present value as in the over-layered strip calculation above.

Mark-to-Market Value

To calculate the discount or premium of a risk-free amortizing bond at a given time during its term, we have to redo the present value calculation and take the difference between the calculation made at the contractual rate and the equivalent rate based on the new yield-curve – for the remaining term of the bond, in both cases discounted using the new yield-curve.

Figure below shows the interest rate for different durations between 2002 and 2009. The 'compression' effect seen from late 2005 to late 2006 corresponds with a flattening of the yield curve.
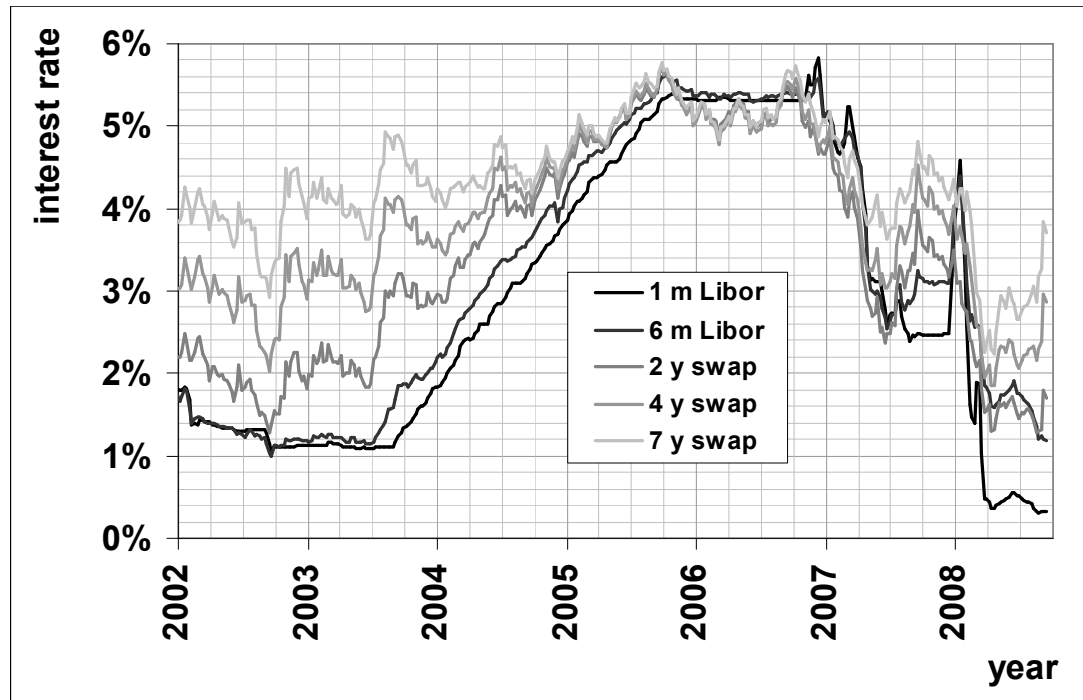
Figure 30: Interest rates for different durations 2002-2009.

If we apply this interest rate evolution on the amortizing exposure above, we get a mark-to-market value that differs from the nominal value of the bond. As shown in Figure , initially, the mark-to-market value was identical to the nominal value. As interest rates rose, the discounting of the remaining cash flows gave a smaller present value than if the discounting had been done at the contractual rate, and the mark-to-market value fell below the nominal value. At the end, there will be fewer and fewer cash flows remaining, and the mark-to-market will converge with the nominal value again.
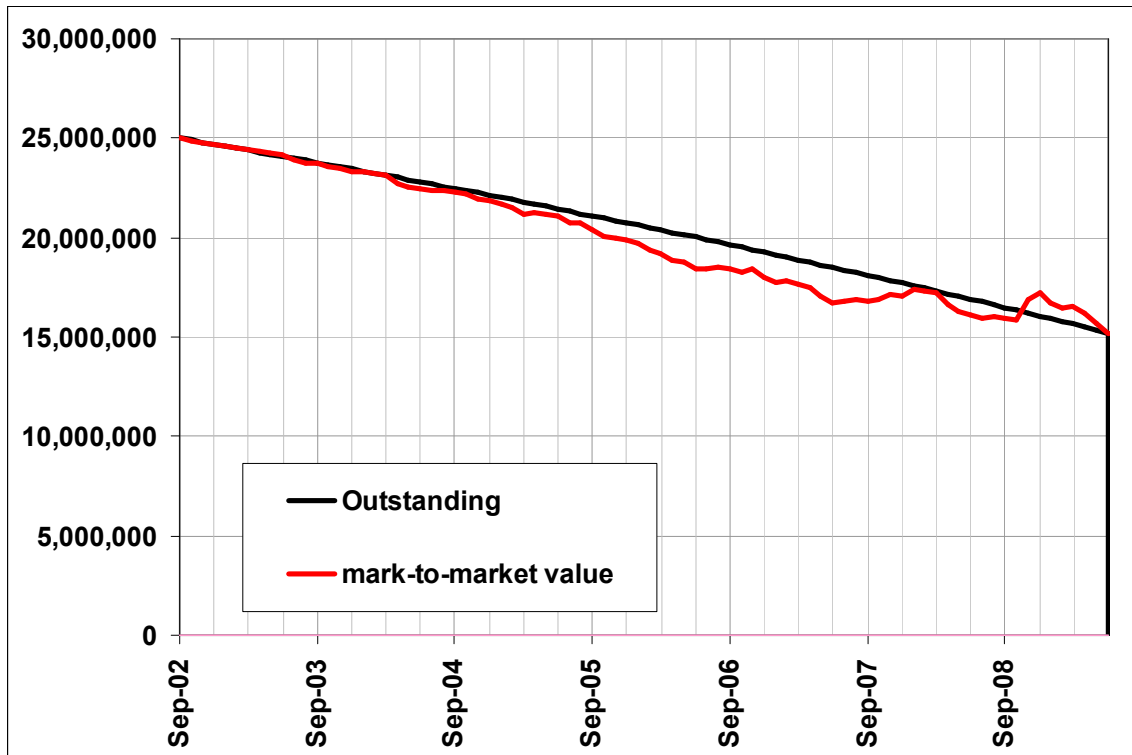
Figure 31: Mark-to-market value versus nominal value of a fixed rate bond when the yield curve evolves.

Forecasting the Yield-Curve

Figure illustrates that market interest rates fluctuate. In this case, we show 1 month Libor.
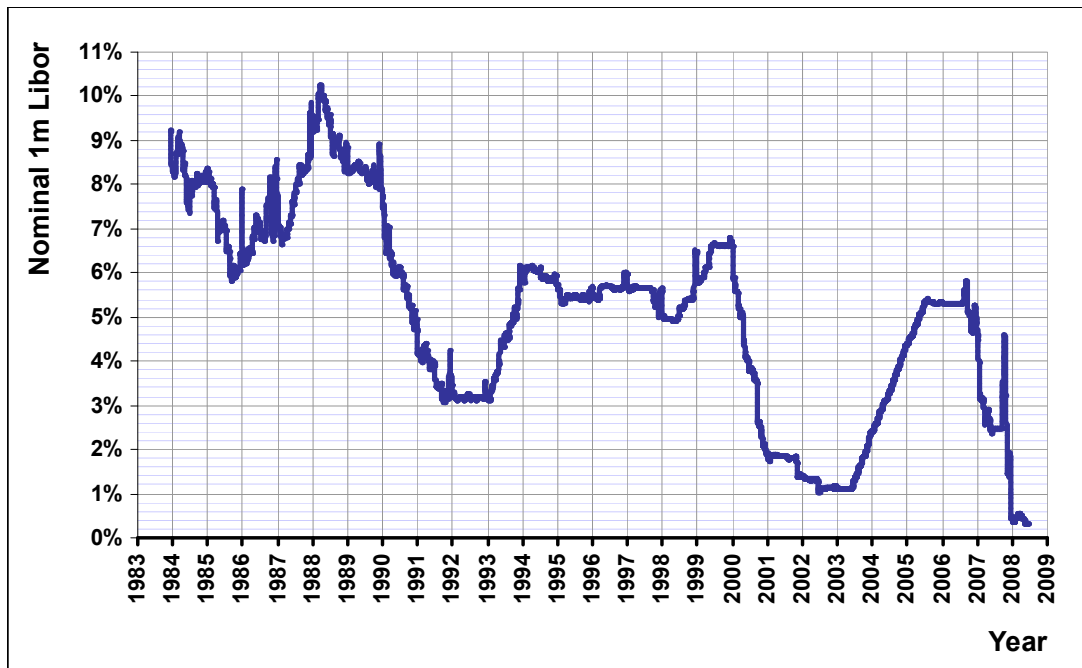
Figure 32: Historical 1m Libor rate since 1984.

In Figure below, we plot the interest rate for different durations, one month, three months, 6 months, one year, and 2, 3, 5, 7, and 10 years, respectively. On this chart, vertical separation between different durations is an indication that the yield-curve is sloping. Note how the yield curve seems to slope more when short term rates are low.
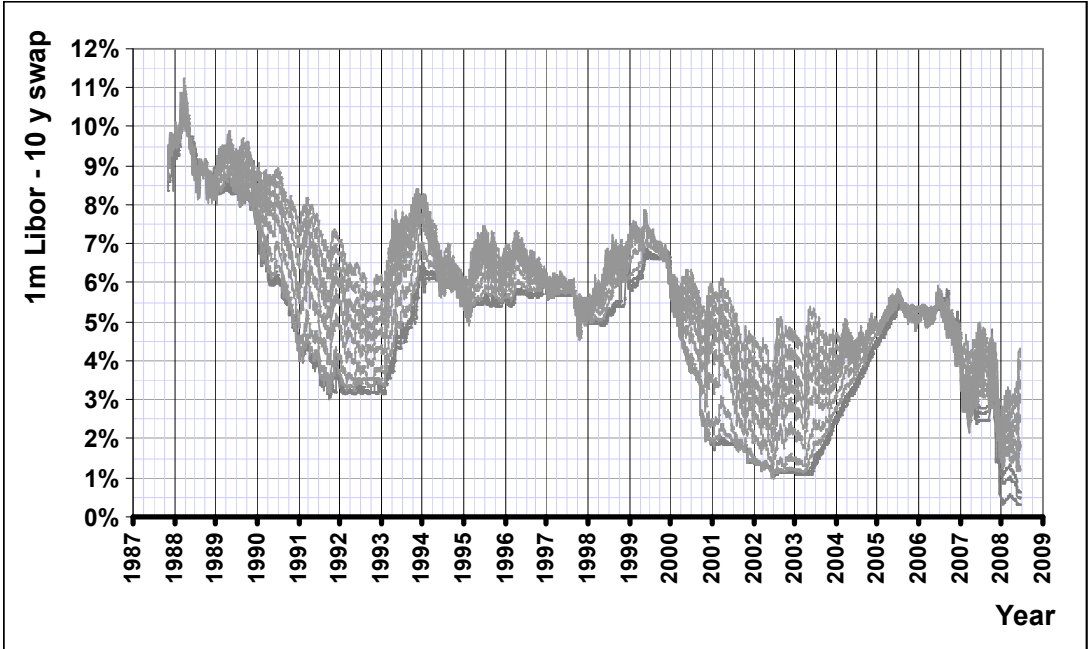
Figure 33: Historical interest rates for durations from 1m Libor to 10y swaps.

An alternative way to illustrate the yield curves for different durations is by showing a "yield surface" over time. Figure below shows rates from 1987 to present in the form of a surface.
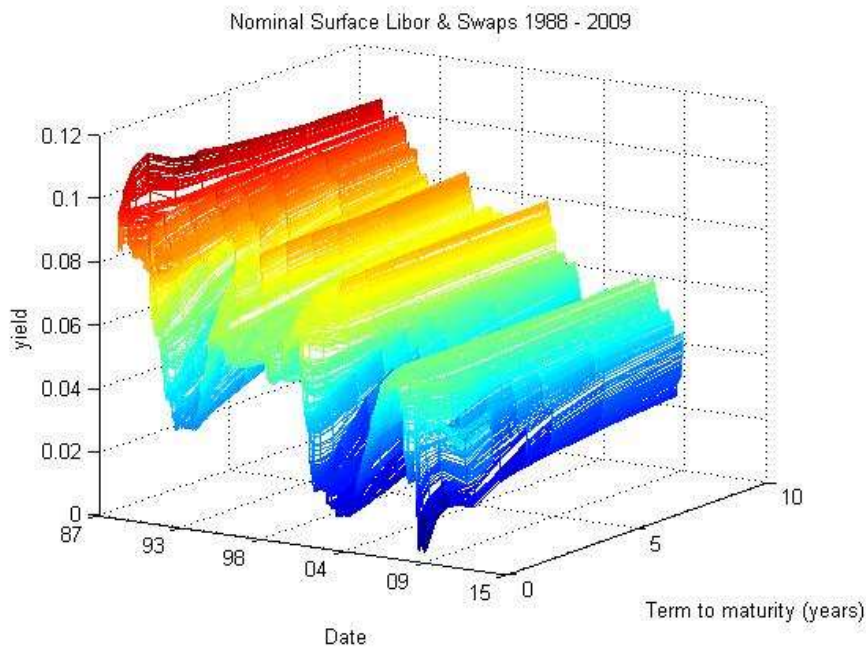
Figure 34: Historical yield surface

There are many different models to simulate the evolution of interest rates. When future interest rates are modeled for purposes of pricing interest rate derivatives (caps, floors, collars, swaptions, etc.), the rates are modeled in the 'risk-neutral' world. In the context of aircraft loans and leases—given the long term and the fact that the interest risk is linked to default risk, and thus un-hedgeable—we will model the yield curve in the 'real' world.

In the illustrations below, we have used the Nelson-Siegel method to interpolate yield curves using three components (basically level, slope and curvature of the yield curve) and a mean-reverting stochastic process (an Ornstein-Uhlenbeck process with inertia) to simulate the yield surface evolution over time. Figure below shows an example of one simulation of the yield curves for durations between two weeks and 10 years. In this example, the short term rate rises, and the slope flattens, and even goes inverted, near the peak of the short term rate. This behavior can be observed in historical data.

Figure 35: One simulation of the yield surface using an Ornstein-Uhlenbeck mean reverting process. The lines represent different durations (here expressed in years).

The Value of a Risk-Free Amortizing Bond

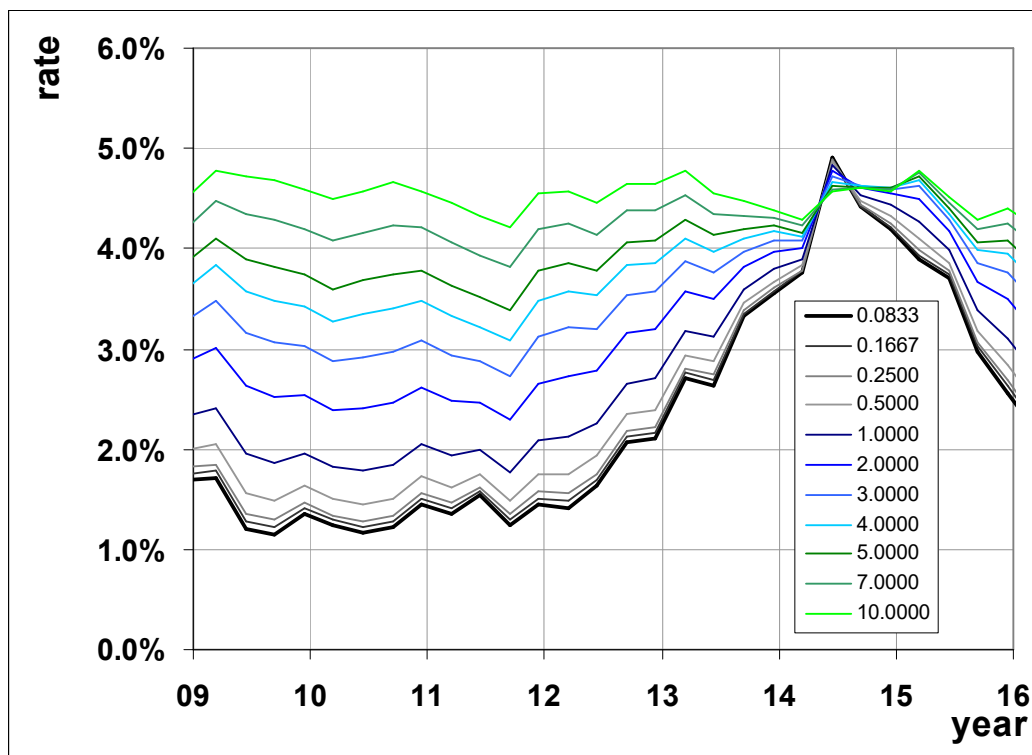Once we have simulated possible evolutions of the yield-curve into the future, we can then calculate the mark-to-market value of the amortizing bond over the term for each evolution of the yield-curve.

Interest Rate Risk and Hedging

When we grant an aircraft loan or invest in an aircraft on operating lease, we have to finance this asset. If the asset yields a fixed interest rate, we are running a risk if we finance the asset by borrowing at a floating rate. An increasing floating rate will increase our borrowing cost while the fixed rate asset will yield the same rate as before. Therefore, it is typical that we fix the borrowing rate with an interest rate swap.

If the lender sells the loan or the lessor sells the aircraft with the lease attached at some point during the term, the mark-to-market adjustment due to interest rate movements will be exactly off-set by the mark-to-market of the swap. Provided the swap counterparty is default free, we will thus have hedged the interest rate risk away. However, if the borrower or lessee is at risk of default, we would carry the risk that the cash flow stream will no longer match the contractual payments that are mirrored in the swap. In this instance, it is the *swap agreement* that will carry a discount or premium relative to the notional amount—the swap breakage cost/gain.

Economically, it doesn't matter if the swap is broken and the swap breakage cost is paid, or if we let the swap carry on. The opportunity cost for keeping the swap is identical to the breakage outcome. In loans, it is typical that the cost of breaking the swap will be due under the loan and covered by the security over the aircraft. This is referred to as swap break cost protection or "make whole premium".

These amounts are also typically due in the case of voluntary pre-payment. The risk here is that the fixed interest rate for the remaining term based on the market yield-curve becomes lower than the fixed contractual rate; the swap will cost money to break. Hence, when we model an aircraft loan or lease, the latent swap breakage has to be added to the exposure.

In Figure below, we show an example of the swap breakage cost outcome under a series of yield surface simulations for a partially amortizing exposure (the one described in Figure 29).
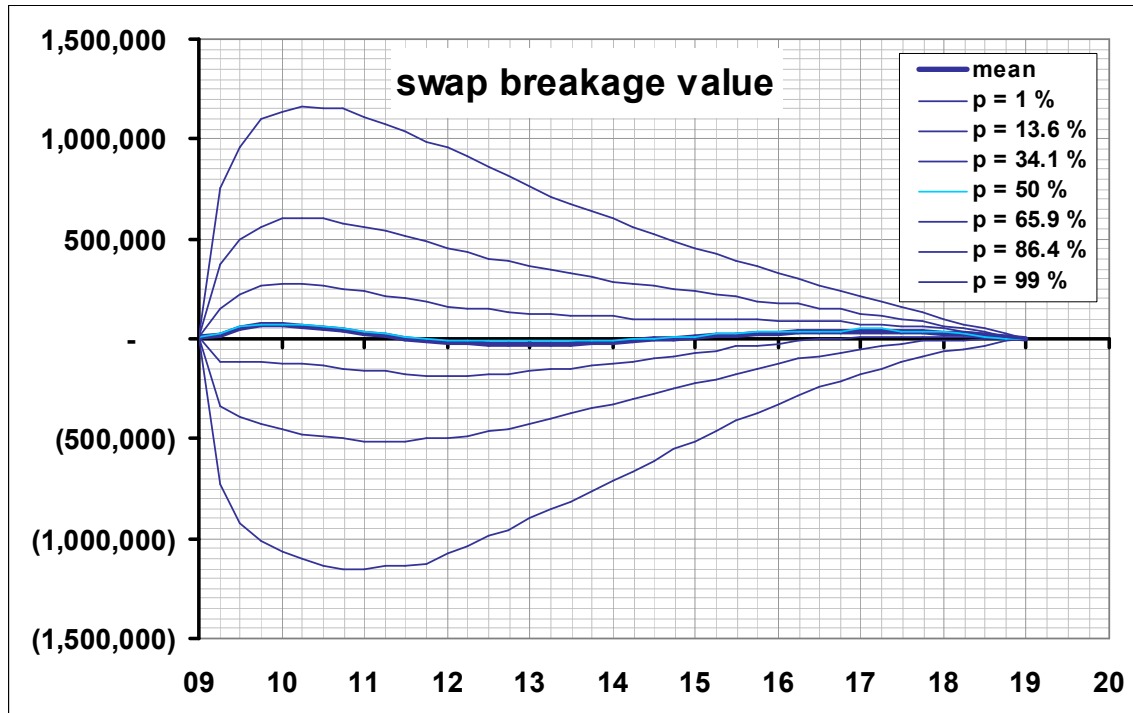
Figure 36: Percentiles of swap break outcomes for an amortizing exposure in a 5000 scenarios simulation of the yield surface.

Floating Rate

In a floating rate transaction, the interest rate risk is limited to the fluctuation in rates over a payment period and is covered under "re-employment of funds" provisions. Given the shorter period, these amounts are often negligible and hardly worth modeling. There is possibly an element of added credit risk in floating rate loans as the periodical debt service could increase with increasing interest rates to a point where the borrower finds it difficult to meet its obligations.

Inflation and Interest Rates

We discussed in the section under Aircraft Values that inflation affects aircraft values over time. Inflation since 1983 is shown in 37 below.

Figure37: USD Consumer Price Index since 1983.

Inflation is correlated to interest rates. We can see the correlation in the scatter diagram in Figure below. This diagram shows monthly US CPI (Consumer Price Index) on the vertical axis against one month Libor on the horizontal axis.

Figure 38: US CPI Inflation plotted against one month Libor (1984-2009)

If the nominal interest rate is adjusted for inflation as per the formula below, we get the **real** interest rate.
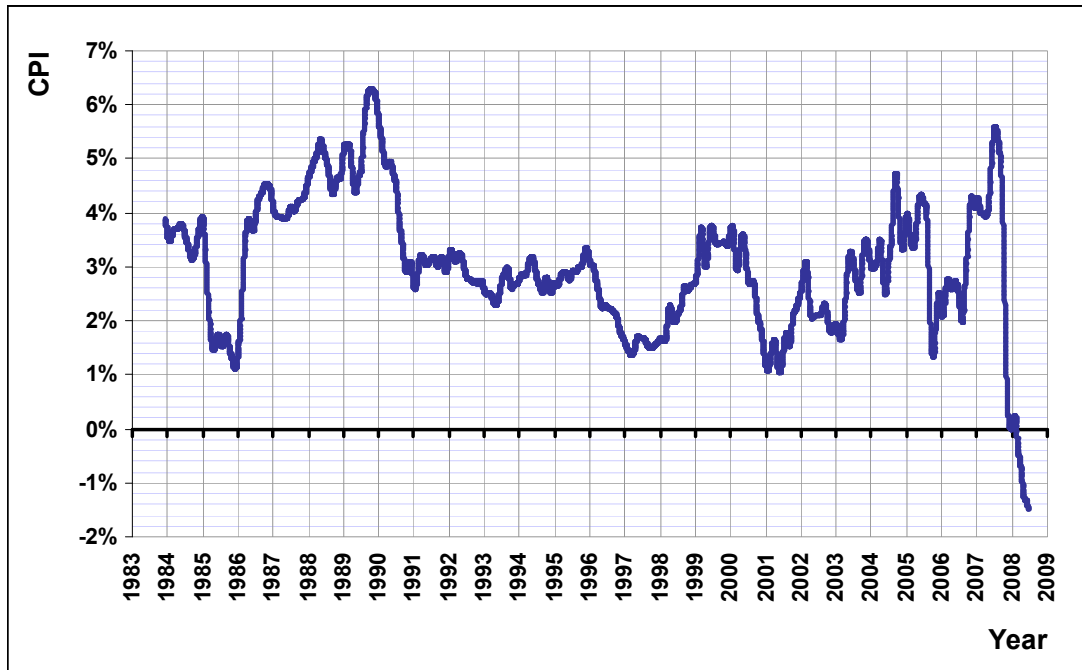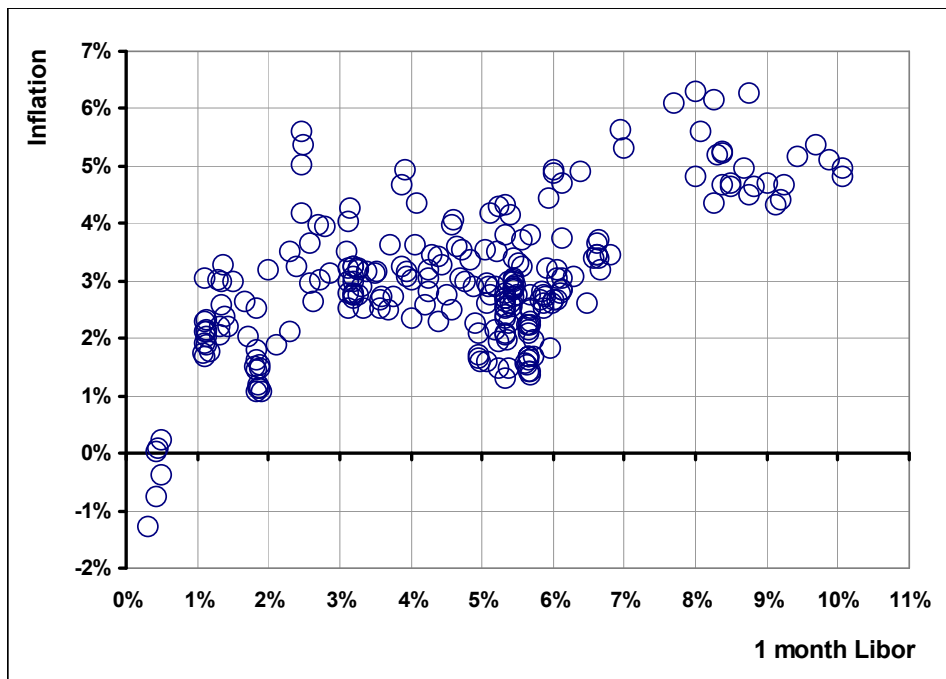
$$(1 + i_n) = (1 + i_r) \times (1 + \pi) \Rightarrow i_r = \frac{(1 + i_n)}{(1 + \pi)} - 1 \quad \text{where } \pi \text{ is the inflation.}$$

Inflation is very difficult to forecast as it depends on economic, monetary and fiscal policies as well as expectations among the public. Some countries have experienced hyperinflation with inflation rates running at hundreds, or even thousands percent per annum. The European Central Bank states as its main objective to keep inflation in check. For the last 20 years, inflation in the USD and Euro zone has been modest. Can we assume that this will be the case for the next 20 years?

Inflation expectations could be revealed in the spread between the yield on Treasury Inflation-Protected Securities (TIPS) and regular Treasury bonds for similar maturities. These securities have been sold since 1997 and now represent about 10% of the total outstanding Treasury debt. Although a liquidity risk premium may blur a direct translation to a market inflation expectation, the TIPS bond can nevertheless be helpful when forecasting inflation.

We could simulate inflation using a mean-reverting random walk and we could tie the mean reversion level in the interest rate simulation to the inflation rate in each scenario. The inflation scenario that we simulate should also be used for the inflation assumption in each simulation of the aircraft value. This is an area that needs more research, but it is an important one, especially when leasing aircraft.

The Cycle

Evidence of cycles

Anybody who has been involved in aircraft financing over a long time will have noticed that airline defaults happen in clusters, that aircraft values swing up and down, and that interest rates move with the economy. There is ample evidence of this cyclicality. Worse, it seems that when airlines tend to default and we rely on the aircraft value to recover our book exposure, it is precisely when aircraft values will be depressed, and interest rates will have fallen to a point where we face swap breakage costs on fixed rate transactions. Therefore, we cannot look at default risk, aircraft value volatility and interest rate volatility in isolation and assume that these are independent variables. In fact, they are highly correlated.

Figure shows the appraised value for a 1988 Boeing B737-300 **aircraft** by Ascend (blue line). The red line is a trendline (a cubic spline) line capturing the long-term depreciation of the aircraft.
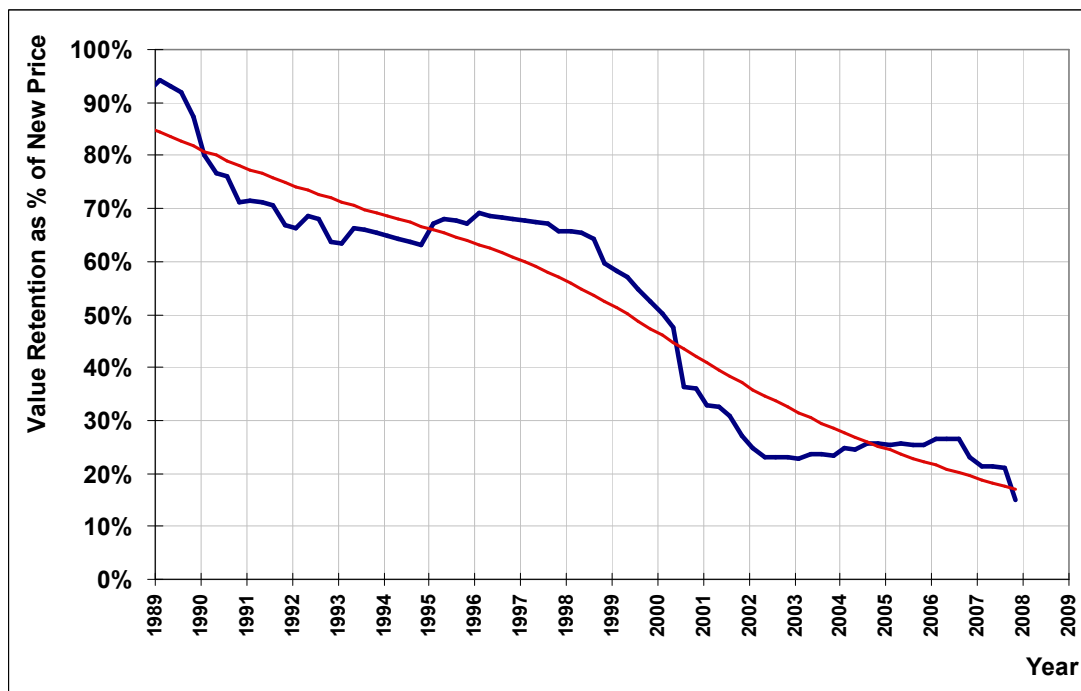
Figure39: Appraised value for 1988 Boeing B737-300 as per Ascend (blue line) with trend line (red line).

If we look at the blue current market value line as a percentage of the red trend-line, we capture the cyclical swings in value. Figure shows this cyclicality. It is easy to identify the 1991-93 "Gulf War I" downturn and the 2001-2003 "September 11" downturn. In 2008, values were on their way down in the "sub-prime" downturn. Note how the amplitude of the swings increased from 10% to 30% over the 20 years the values were measured.
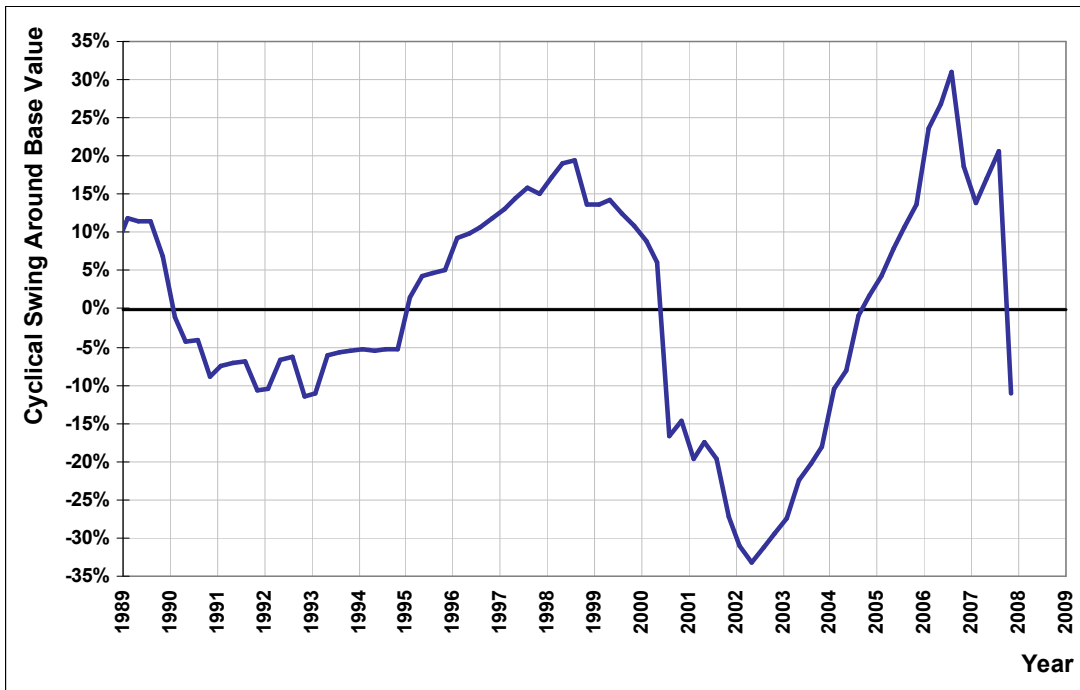
Figure 40: Relative cyclical value swings of a 1988 B737-300.

Obligor credit worthiness also displays a cyclical behaviour as evidenced in Figure . The median, upper and lower quartile EDFs for airlines as per Moody's KMV are shown. Higher EDFs (weaker credit quality) are evident in 1991, 2002, and 2008.
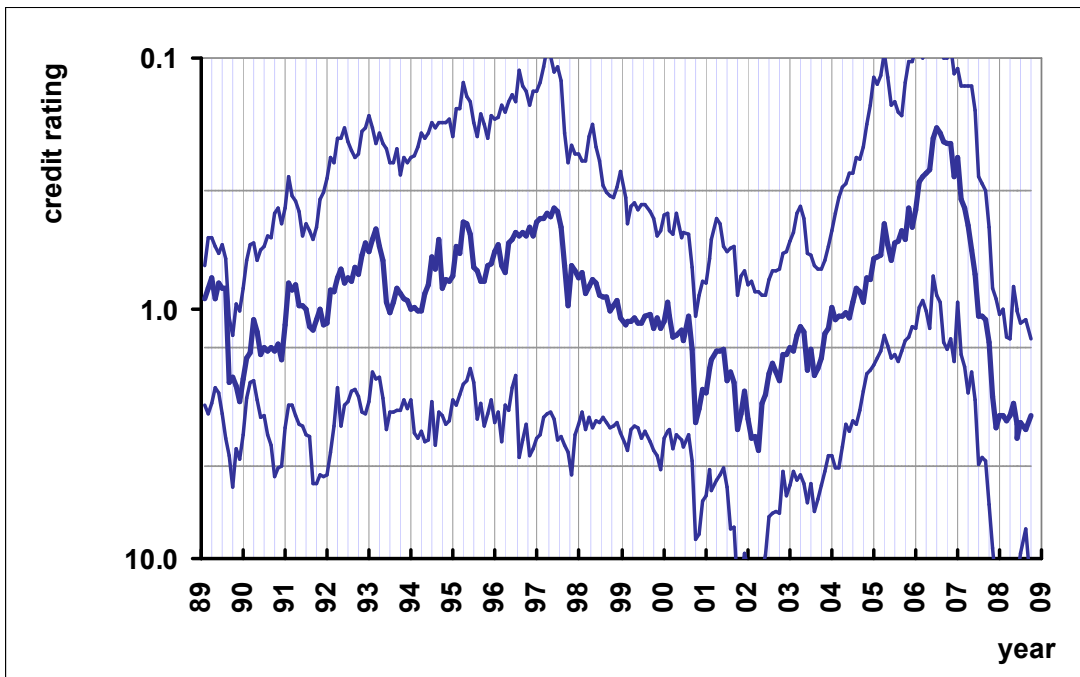
Figure 41: Airline credit rating median and quartiles over times using Moody's KMV EDFs. The cyclicality is evident.

Interest rates are also cyclical. The time series for one month Libor is shown in Figure below. Although the trend has been falling, the rate has still displayed a cyclical pattern. Again, we see troughs in 1992, 2002, and 2008.

51
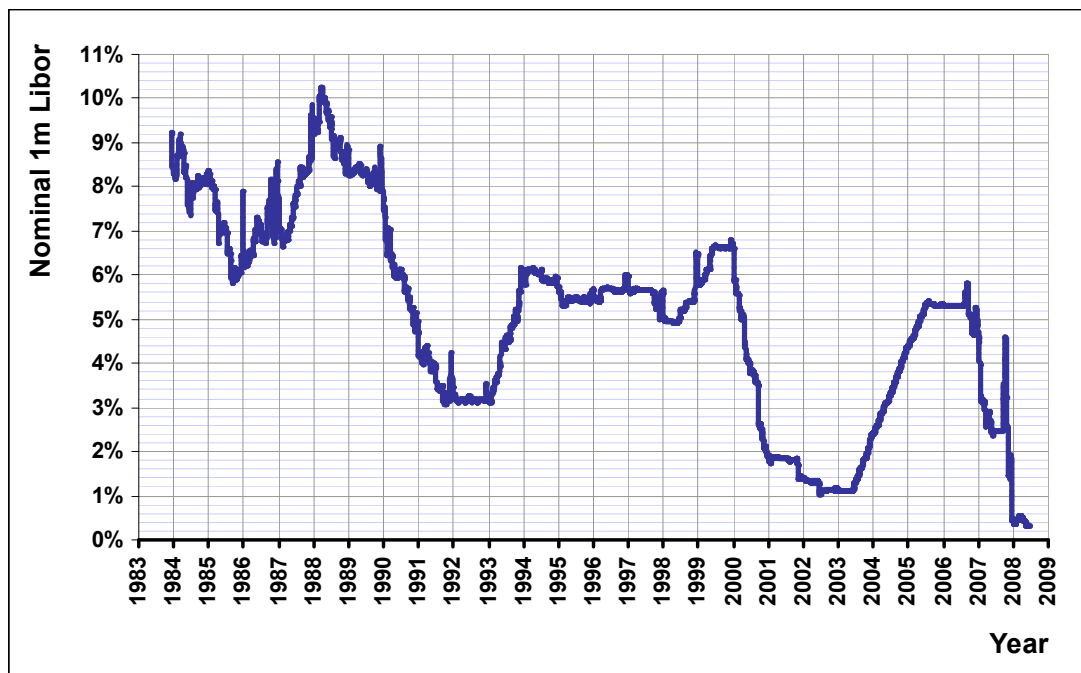
Figure 42: One month Libor rates.

Correlation and Causation

It is interesting to note that, at least over the limited time period we have observed, the cyclical troughs and peaks of aircraft values, airline credit worthiness and interest rates seem to coincide. Could it be that there is a common cause that would explain this?

The traditional approach to deal with dependent variables in valuation and risk models is to measure the historical correlation and use this correlation to model the future. We would thus assure that when the aircraft value is simulated to be comparatively low in one scenario, the default probability would be correspondingly higher, and maybe the interest rate lower in the same scenario. And vice versa.

Another way is to see if there could be a common cause for the observed correlation and try to simulate this going forward. The common cause would affect default probability, aircraft value and interest rates simultaneously in each scenario. Even better, if we thought we had some ability to predict the common cause, we could use that in the simulation.

It is intuitively compelling to think that aircraft values would hold up when there is strong demand and fall when there is a surplus in capacity.

Demand for aircraft is driven by air traffic (passenger seat miles, RPM) and capacity can be measured in available seats (number of flyable aircraft times the number of seats installed). So high traffic would cause airlines to buy more aircraft, and when there is a shortage, prices would rise. But in this scenario, airlines would be able to sell a lot of tickets and keep ticket prices high, so they would make a lot of money and the risk of default would be low. Strong traffic goes hand in hand with a strong economy, and when the economy is strong, interest rates tend to be higher.

Let us measure the historical annual traffic growth and seat growth (delivery of new seats less scrapping). Figure shows an extract from the data set (with the projections blurred out).

YEAR	RPM GROWTH	SEAT GROWTH	GROWTH DELTA	ACC DELTA	TREND ADJUSTM	CYCLE
1990	7.00%	7.64%	-0.64%	-70.25%	-75.88%	5.64%
1991	-3.28%	9.57%	-12.85%	-83.10%	-75.19%	-7.91%
1992	12.93%	7.90%	5.03%	-78.07%	-74.47%	-3.60%
1993	2.45%	5.23%	-2.78%	-80.85%	-73.71%	-7.14%
1994	8.48%	2.99%	5.49%	-75.36%	-72.93%	-2.43%
1995	11.09%	2.21%	8.88%	-66.48%	-72.11%	5.63%
1996	8.06%	2.76%	5.30%	-61.18%	-71.27%	10.09%
1997	4.62%	4.21%	0.41%	-60.77%	-70.40%	9.63%
1998	2.66%	5.39%	-2.73%	-63.50%	-69.51%	6.01%
1999	5.35%	6.18%	-0.83%	-64.33%	-68.59%	4.25%
2000	8.63%	5.25%	3.39%	-60.95%	-67.64%	6.69%
2001	-2.55%	5.11%	-7.66%	-68.61%	-66.68%	-1.93%
2002	-0.20%	3.68%	-3.88%	-72.49%	-65.69%	-6.80%
2003	1.88%	2.96%	-1.08%	-73.56%	-64.67%	-8.89%
2004	12.64%	3.29%	9.35%	-64.21%	-63.64%	-0.58%
2005	7.00%	2.92%	4.08%	-60.13%	-62.59%	2.45%
2006	5.00%	3.47%	1.53%	-58.60%	-61.50%	2.90%
2007	5.00%	3.30%	2.90%	-56.10%	-60.41%	4.31%

Figure 43: The annual traffic growth and annual physical seat growth is recorded for past years and forecasted for future years and the delta is calculated.

Let us now take the difference between the annual traffic growth and the annual seat growth, and accumulate this difference over time.
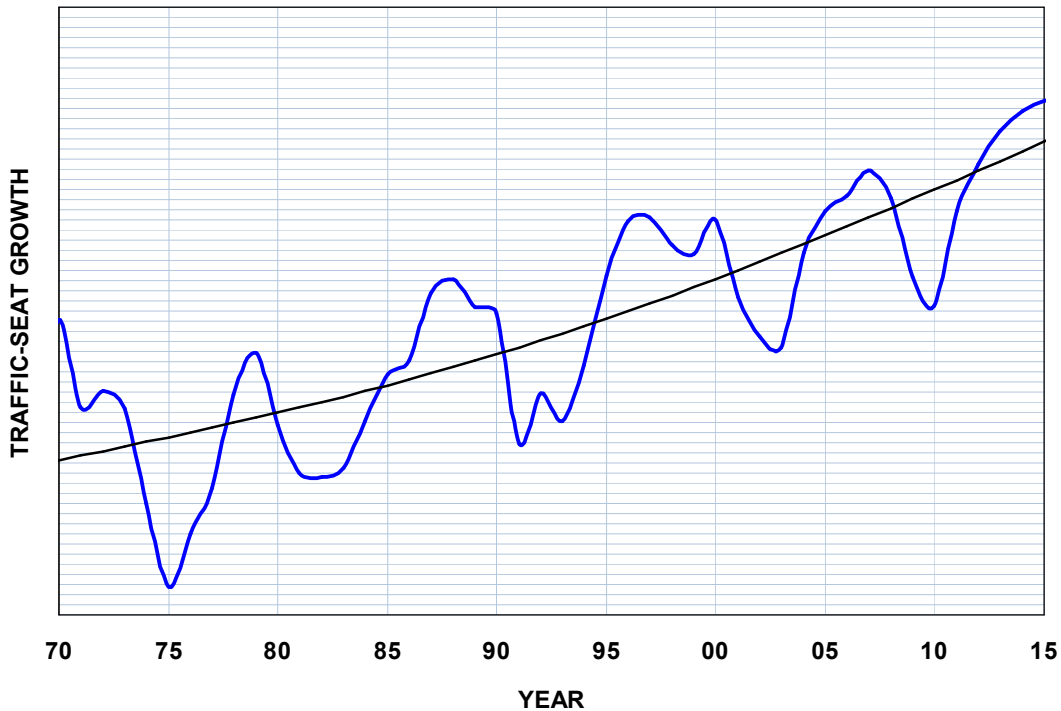
Figure 44: The accumulated delta between traffic growth and seat growth is plotted in the blue line. The long term trend line is calculated (again using a cubic spline).

The growing trend is an indication that aircraft seats have become more productive over time, from aircraft flying faster when production shifted from propeller aircraft to jet aircraft, and then from higher load factors and higher daily utilization of aircraft. This trend is a long term

one, and the airlines and aircraft manufacturers have time to adapt to this. The remaining movement is the cyclical component seen in Figure below. The line can be seen as the Pent-Up Relative Capacity Shortage (positive values) or Surplus (negative values) (PURCS) of aircraft compared to an average market.
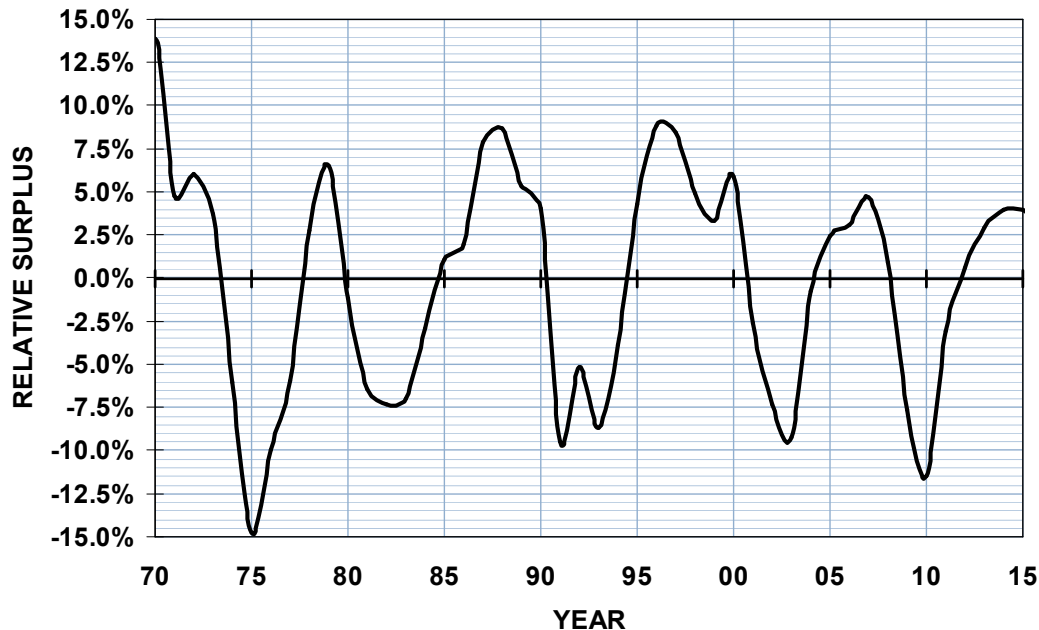

Figure 45: The cyclical component is isolated from the graph in Figure 44. We refer to this as the Pent-Up Relative Capacity Surplus/Shortage (PURCS) cycle.

We see that the trough and peak points in this PURCS cycle seem to coincide with the trough and peak points in the historical time series of B737-300 values, airline EDFs and one month Libor rates. We can look at other aircraft types and see that the correlation to PURCS has persisted since the beginning of the jet age. We have not found airline credit quality or default data dating prior to 1989, but regulation may well have kept the correlation to PURCS weak. Interest rate time series that pre-date the Libor series we showed are not visibly correlated to PURCS. We will discuss this later on in this article.

Forecasting the Cycle

We have seen that aircraft values, default frequency and interest rates all seem to be affected by the cycle. If we were able forecast the cycle, it would help us improve our forecast of aircraft values, default frequency and interest rates.

There are several reasons to suspect that the cycle is somewhat predictable. Just as a pendulum swings around an equilibrium point and gravity tries to pull it back when it is deflected, the capacity cycle works the same way. When there are too many aircraft, demand will fall and the manufacturers will slow the production. Over time, the excess capacity will be absorbed and with the lack of new deliveries, a shortage will develop.

Now, the airlines are making money and new aircraft orders will be placed. Thus, by knowing if the cycle is low, we know the odds are that it will go higher from there, or vice versa. The manufacturers will respond but it typically takes 18-24 months to build an airplane, so the response will be delayed. With few manufacturers and a significant lead time, the deliveries in

the next one to three years are predictable with a meaningful degree of accuracy. Although deliveries are only half of the capacity equation, the fact that they are predictable with a meaningful but declining degree of accuracy will assure some degree of accuracy in the prediction of the cycle itself.

We'll discuss one model for simulating the cycle here. We assume that the cycle consists of an upswing with a wavelength T1 and a trough to peak amplitude of A1, followed by a downswing with wavelength T2 and amplitude A2, and so on. The segments are assumed to be shaped like a sine wave. We can measure the average and standard deviation of the upswing duration and downswing duration as well as the average and standard deviation of the amplitude on the historical data set. It turns out that the downswings are a bit shorter than the upswings (the boom-bust syndrome). With this simple model, and by fixing the point of the most recent peak or trough point, we can simulate scenarios for the future cycle.
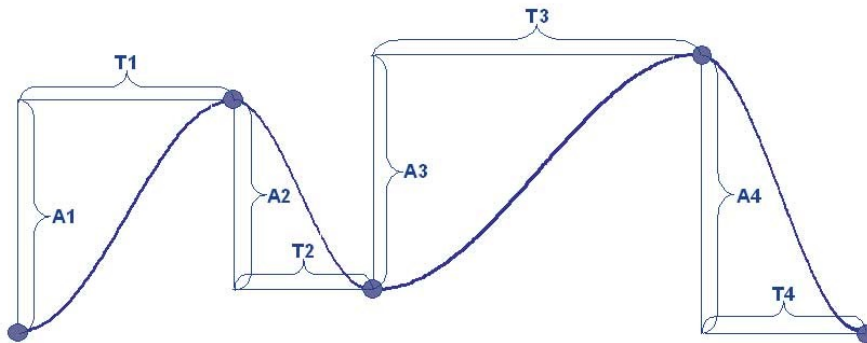
Figure 46: The cycle is simulated with upswing segments and downswing segments. These have random amplitude and wavelength. This is only one way of simulating a cycle. Other approaches could be to simulate traffic growth and seat growth separately.

One such scenario is shown in Figure below. Note that this is just one of many possible scenarios, but together, the wavelengths and amplitudes would have an average and standard deviation consistent with historical data.
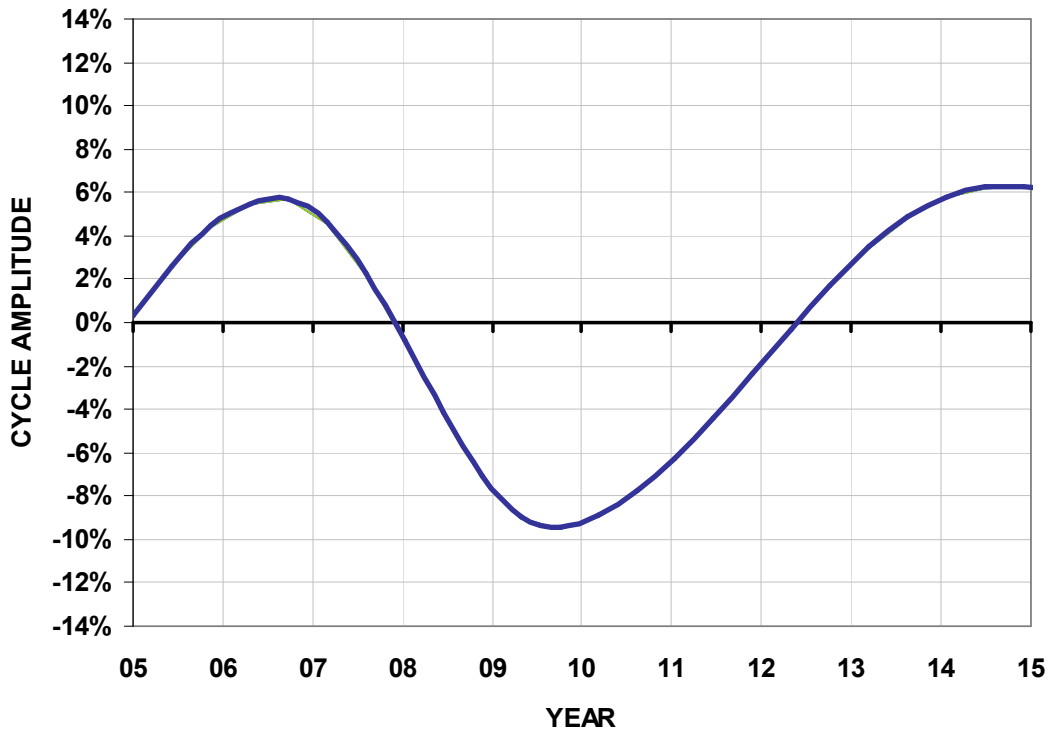
Figure 47: Example of one cycle scenario.

When we simulate 500,000 scenarios, we see in Figure below that the average cycle becomes dampened (like an exponential decay or autoregressive function), as our ability to forecast declines with the increasing horizon. However, each simulated scenario has a full cycle into the future. As the cycle will affect the shape of the simulated aircraft value path, the obligor survival curve and the projected yield curve evolution, we obtain the correlation between these by explicit causation.
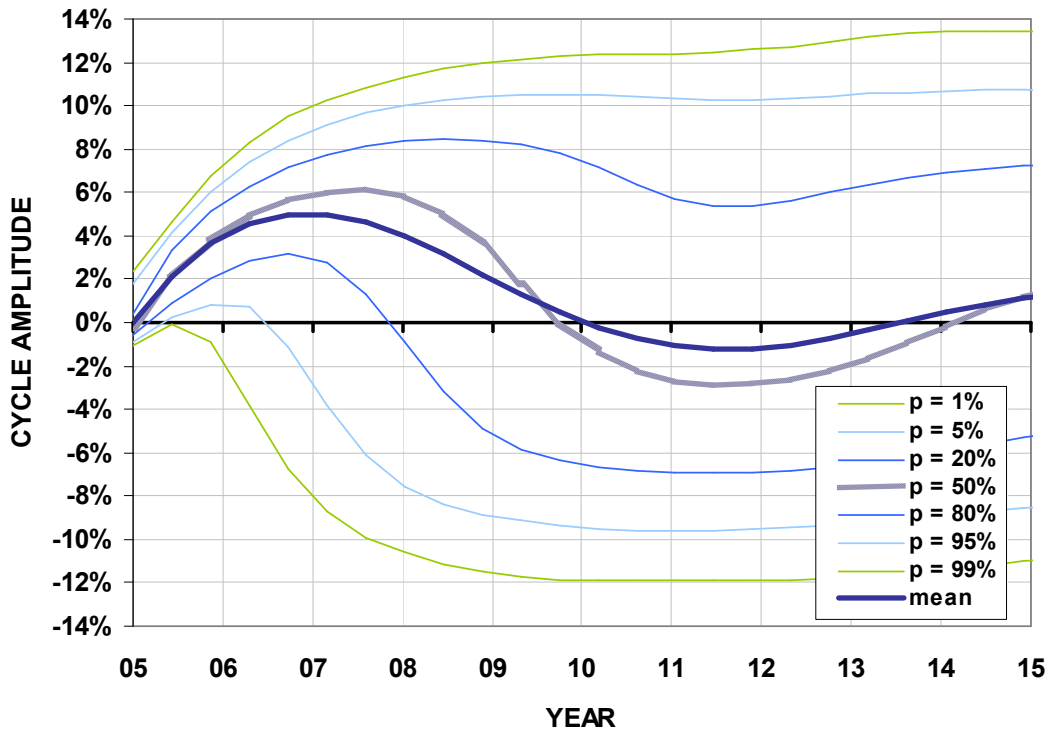
Figure 48: Average, median, and percentiles of 500,000 cycle scenarios. Note how the average is dampened out. With Normally-distributed segment wavelengths, the dampening follows an exponential decay function.

It is not possible to know exactly where we are in the cycle at any given point in time, the empirical cycle will find its final shape only years after the fact. This is due to the fact that we cannot know which part of the capacity growth or traffic growth is due to secular shifts or cyclical shifts. This is why the simulated cycle is subject to variation already in T=0.

PURCS Forecasting Accuracy

The two main components that go into the PURCS cycle are air traffic growth and aircraft seat growth. Ed Greenslet (the editor of the Airline Monitor) has forecasted annual traffic growth and aircraft deliveries since 1984. His forecasting record for air traffic growth is relatively good on average, but it is clear that the forecasting error is not significantly different for different forecasting horizons. Aircraft deliveries, on the other hand, have been predicted with smaller errors for short horizons than for long horizons. This is expected since aircraft take 18-24 months to build. We can tell with a high degree of precision how many aircraft will be delivered over the next 12 months, a little less precision for the following 12 month period, and so on. As a statistic, aircraft deliveries is a reasonable proxy for aircraft seat growth since the seat per aircraft average only changes slowly over time, and scrapping is a result of the cycle, more than a driver of it.

We have recorded our PURCS forecasts since 1995 and compared them to the actual outcome. Figure shows the 14 forecasts (over 10 years) against what we think the actual PURCS cycle looks like. Clearly, the forecast record is very far from being perfect. We don't even know exactly where we are in the cycle at any given point in time. But the record is far better than a random guess. In fact, it seems that the standard deviation of the initial error is about 50% of the average amplitude of the troughs and the peaks, increasing to around 100% after 5 years. This is the volatility we should strive to replicate in our simulations above. Even if our ability

57

to predict the cycle is limited, we can think of it as counting cards in Black Jack – know when the odds are turning in our favor.
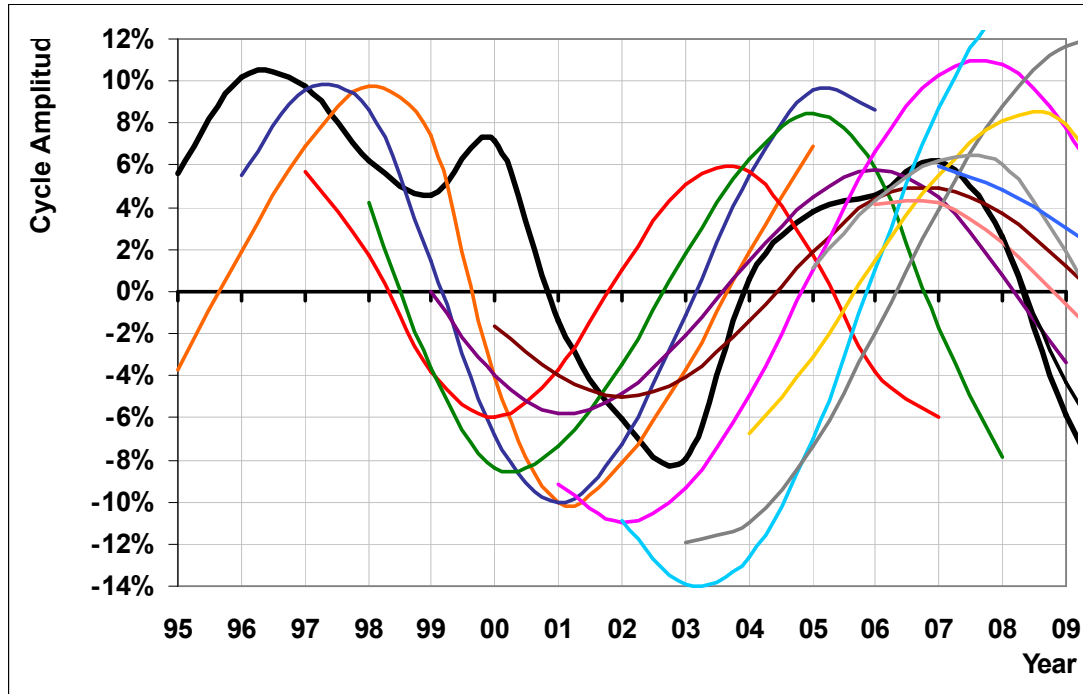
Figure 49: The PURCS cycle (black line) plotted against subsequent 10 years PURCS projections made in 1995, 1996, 1997, ... etc.

PURCS' Effect on Aircraft Values

We have seen that historical aircraft values were highly cyclical. When air traffic growth is strong, many people fly, and when aircraft capacity is short, ticket prices hold strong. This leads to strong demand for aircraft and prices go up. To account for this effect in our aircraft value forecasts, we let the generated PURCS cycle scenarios modify the projected cycle-neutral aircraft value path.

Different aircraft types are affected more or less by the cycle as discussed in the section on Aircraft Value Volatility. The sensitivity to the cycle can be expressed in an aircraft score. The very best aircraft show a sensitivity of around 1.5 – meaning that the aircraft value will swing to 1.5 times the PURCS amplitude from base value to current market value. Old and big aircraft have seen sensitivity factors of 3 or more. Although old and big aircraft become very risky investments, they can also generate the largest gains provided the investment entry and exit are timed right. Towards the end of the economic life of an aircraft type, there will be one last cycle when the value will no longer swing up in an up-cycle. That trough will be the death-trough and it is very difficult to call beforehand.
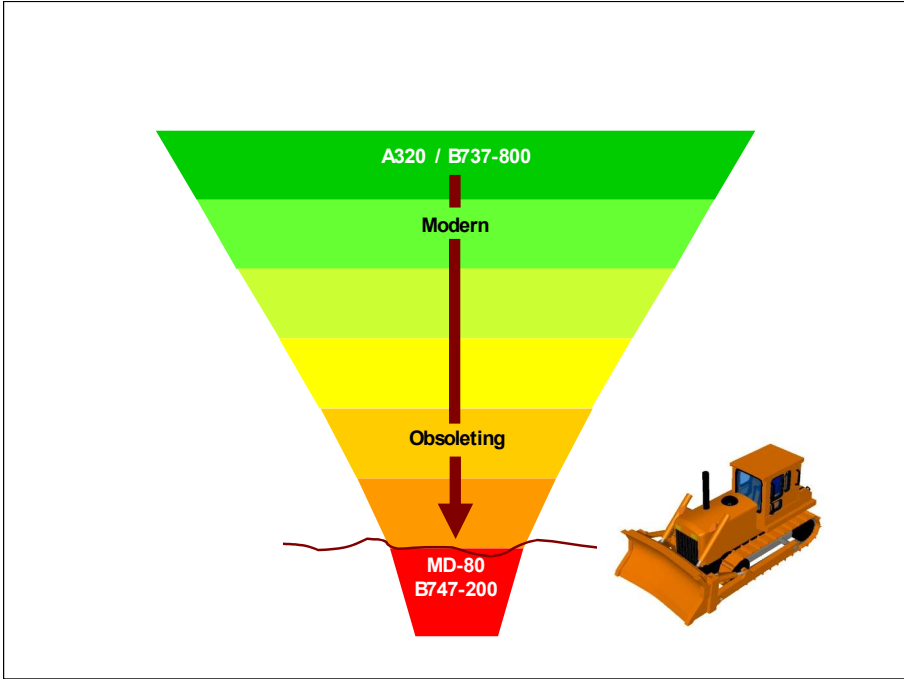
Figure 50: The "aircraft type funnel" illustrates that in every downturn, a portion of the aircraft population will never recover in the ensuing up-swing.

If we go back to our aircraft value simulation from pages 36-39, and add the effect of the cycle scenarios generated as per above, we can simulate the aircraft value taking the cycle into account. One possible path is shown in Figure .
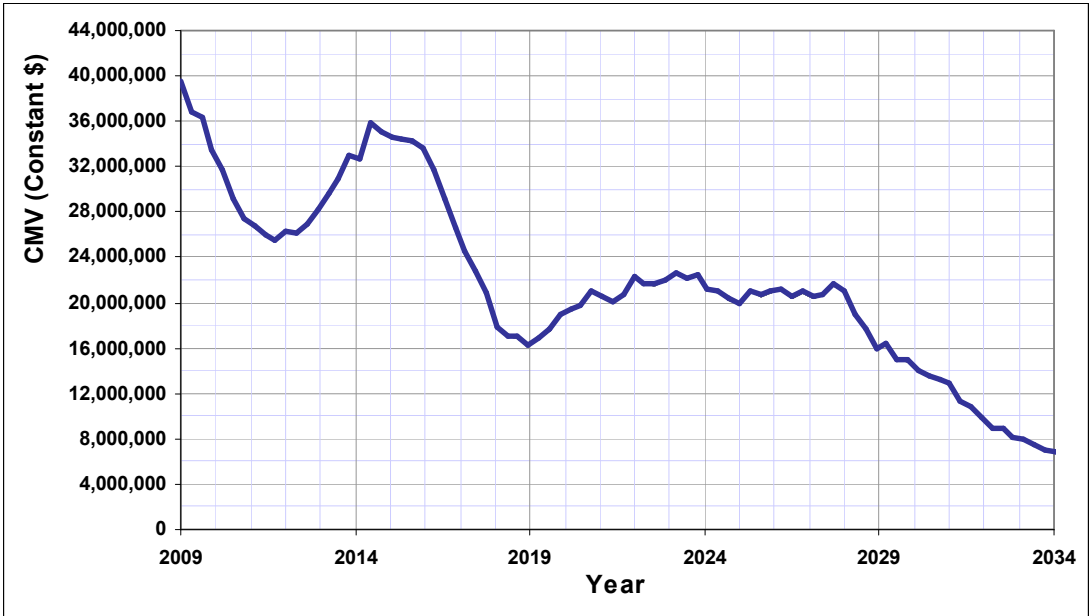
Figure 51: One scenario for the simulated aircraft value adjusted to reflect the cycle.

If we do 500,000 simulations, the result is shown in Figure . The effect of the cycle is visible in the average over the first 5-6 years.
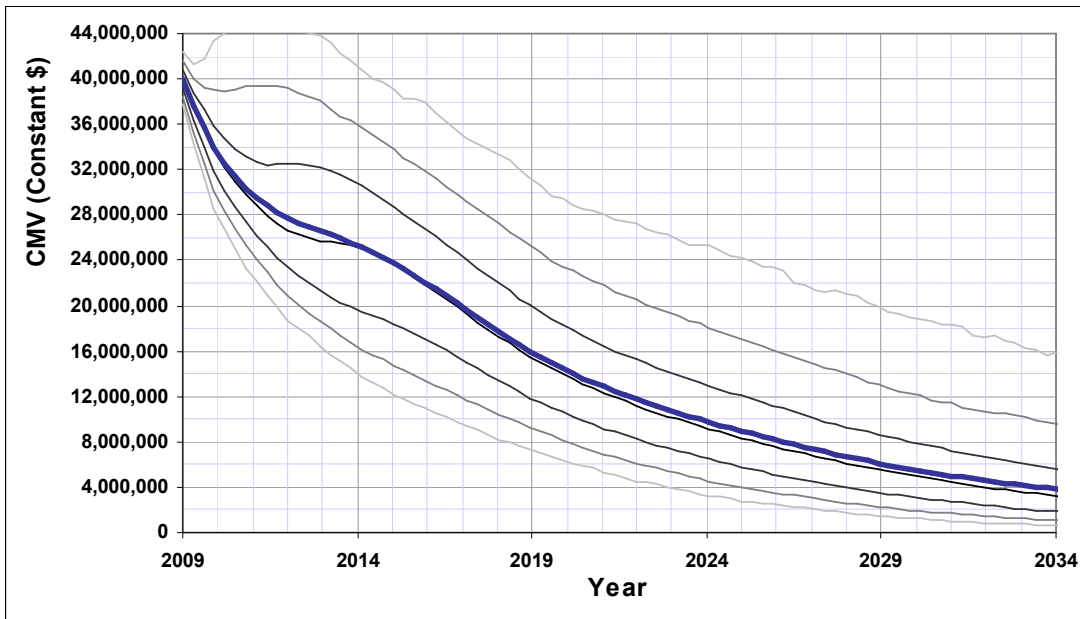
Figure 52: Percentiles of the assumed aircraft value path modified for the existence of a cycle. In this case, it is expected that the cycle will tend down in the early years and rebound towards 2015. But many different scenarios are possible within this expectation.

PURCS' Effect on Obligor Survival

We have seen that the Expected Default Frequency as per Moody's KMV of airlines is cyclical. We also know that airline defaults tend to cluster during periods of weak air traffic and aircraft values. When air traffic is strong and ticket prices are high, the airlines do very well and the risk of running out of cash or becoming insolvent diminishes. One possible way to account for this effect is to use the cycle neutral obligor survival projection and stretch the horizontal time axis in times of peak cycle and compress it in times of trough cycle. It is like exposing the obligor for more stress during a downturn and less stress during an upswing. We have attempted to calibrate this effect based on the KMV data.
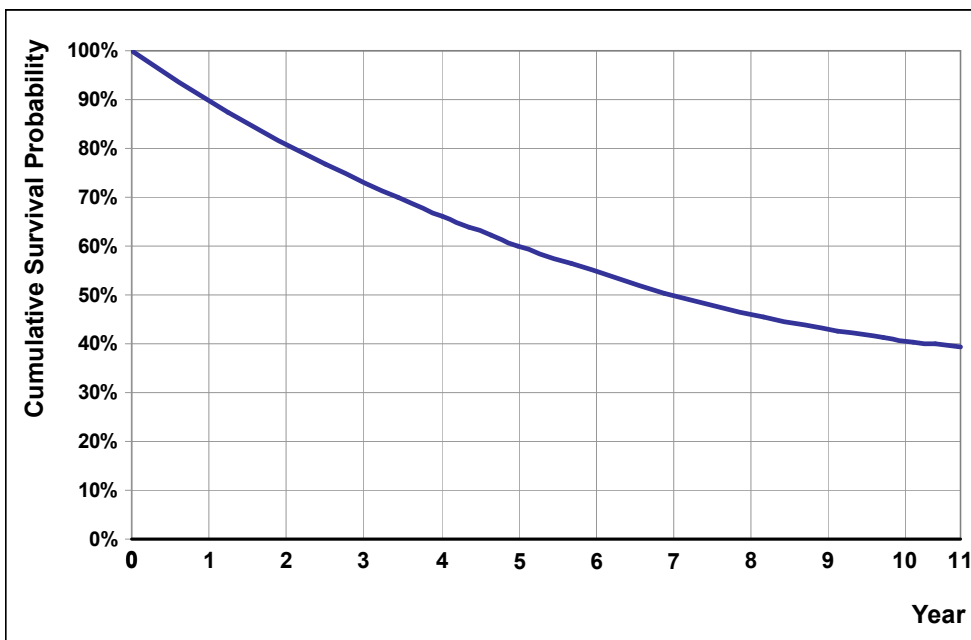
Figure 53: The average survival probability for a specific rating grade based on empirical data collected over many cycles.

60
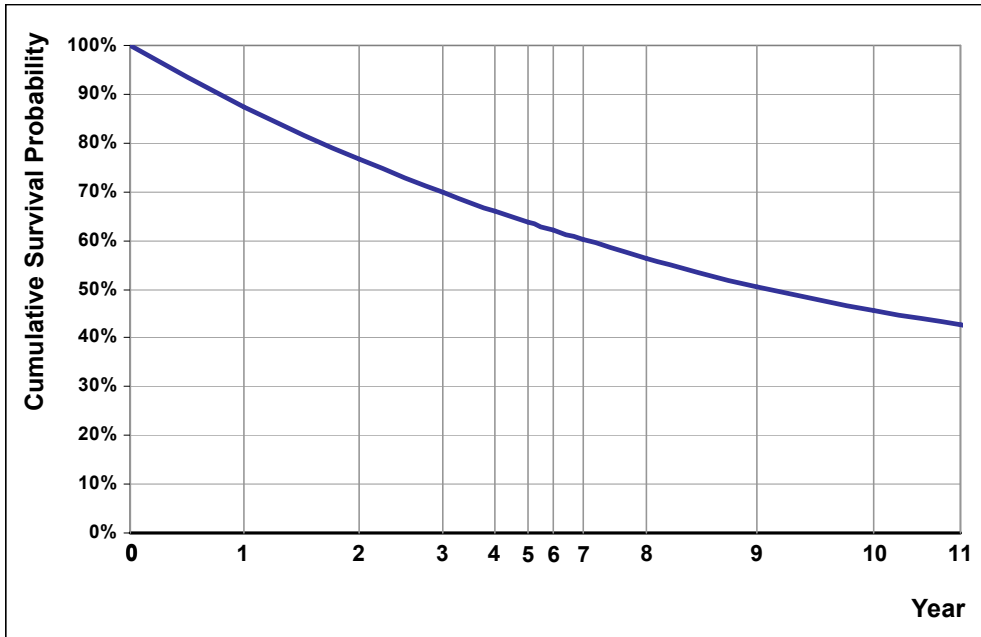
Figure 54: The same survival probability curve as in Figure 53, but with the time axis compressed and expanded to represent passages through down cycles and up cycles.
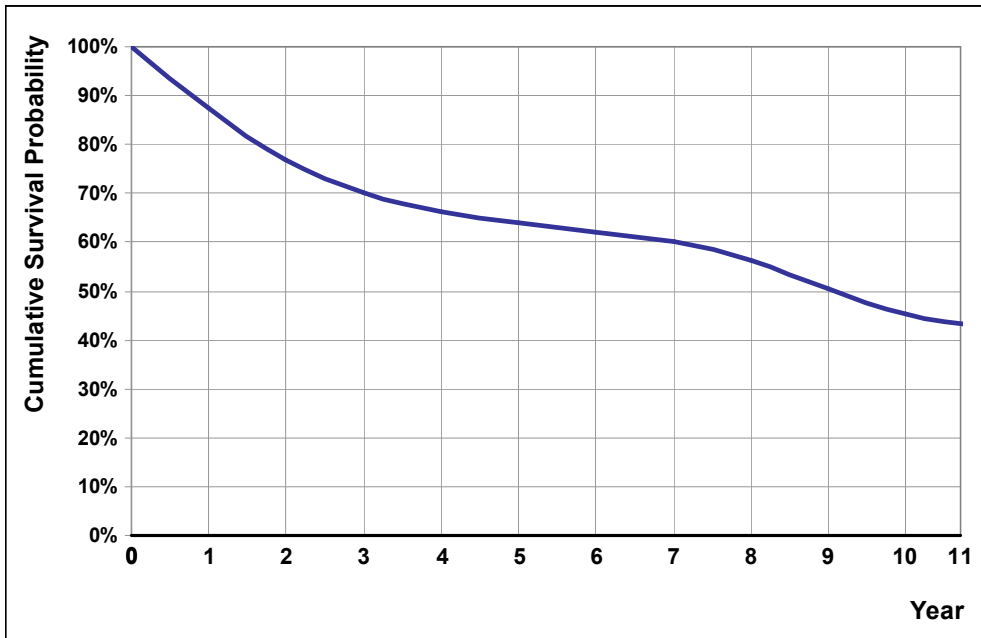
Figure 55: The survival probability curve from Figure 54, but with the time axis adjusted to equidistant steps. Note the lower assumed default probability in an up cycle evidenced by the shallower slope of the survival curve around year 4-7.

By generating a uniformly distributed random variable between 0 and 1, we obtain a time to default as the intersection with the blue line in Figure . Note that any random number below 0.43 would result in no default over the tenor (in this case 11 years) in this particular scenario. Note also that the cycle is simulated for each scenario, so the compression – expansion of the basic survival term structure for the assigned rating grade will be different in each scenario.

PURCS' Effect on Interest Rates

While the PURCS cycle is a very compelling explanation for the cyclical behavior of aircraft values and airline default, it is less clear that interest rates would be correlated to PURCS. One of the two main components in PURCS is the seat capacity growth. There is no strong reason why this would affect interest rates. However, the other component is air traffic growth.

Air traffic is strongly correlated to GDP growth and it is plausible that GDP growth and interest rates are interrelated. So chances are that there is at least a weak correlation between PURCS and interest rates.

One single scenario of the yield surface and the PURCS cycle is shown in Figure . We can see that the short term rate is higher in PURCS peaks and lower in PURCS troughs. We have modified the mean-reverting process described above to include a drift-term that has a weak link to the simulated PURCS cycle.
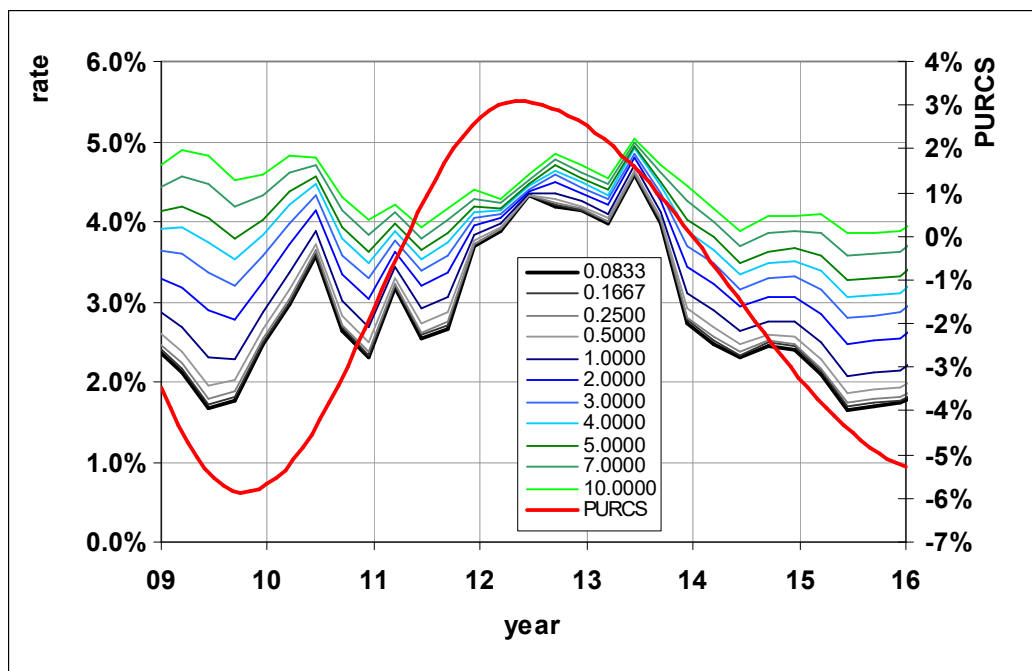
Figure 56: Yield surface scenario with PURCS scenario

If we run 50,000 scenarios of the yield surface simulated with a link to simulated PURCS scenarios, we get the swap break distribution in Figure below.
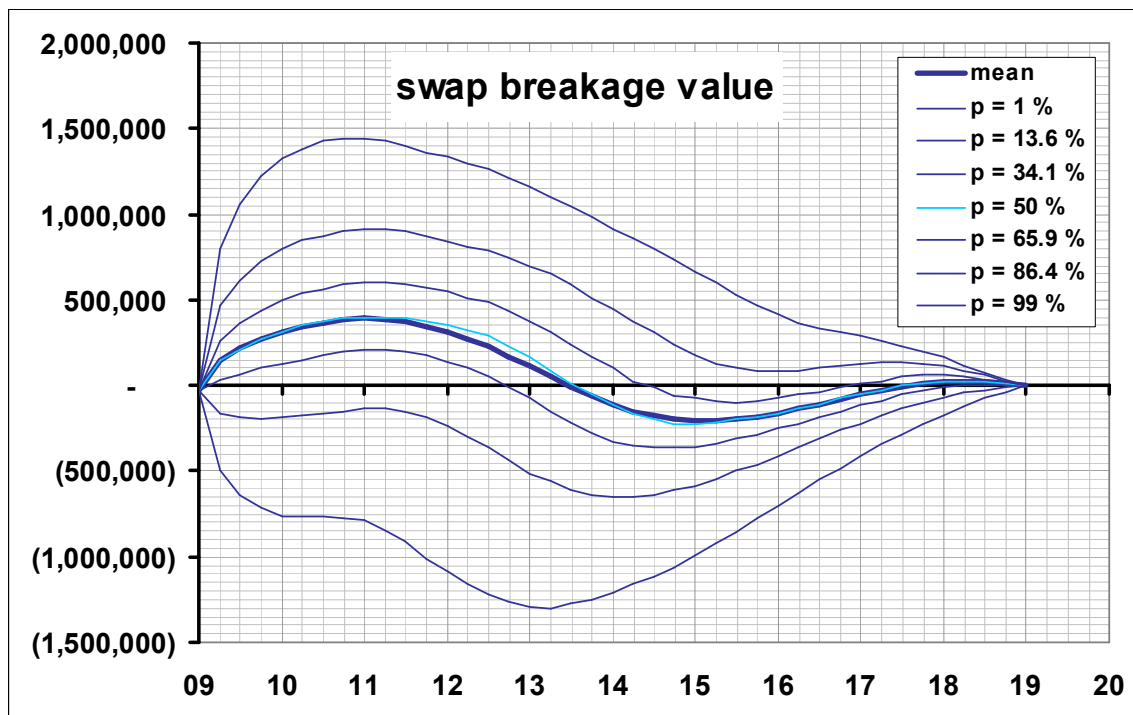
Figure 57: Swap break distribution when the yield surface evolution is influenced by PURCS.

The reader may compare the break value distribution in Figure to the one in Figure where no cycle was assumed.

A Note on Arbitrage

The belief that one could at least partially predict a cycle that would affect aircraft values, default intensity and interest rates, and thereby the value of a loan or lease, would somehow imply that there could be a "house edge", a statistical arbitrage to be exploited. This flies in the face of the no-arbitrage assumption that is a cornerstone in many quantitative financial models. However, the no-arbitrage assumption is mainly relevant for modeling assets in the "risk-neutral" world, in complete markets (see page 82). The market for aircraft loans and leases is incomplete, with few participants, private, opaque, with high transaction costs and low trading frequency. Statistical arbitrages may persist for long periods. And herein lies an opportunity for the participants!

Loan and Lease Pay-off Functions

We have discussed how we can forecast (or guess) the default probability, aircraft value, and the interest rate over the horizon of a contemplated aircraft backed loan or lease. In order to calculate the net present value distribution of the transaction, we need to look at the contractual pay-off under different scenarios.

To conceptualize this, we can think of the transaction as a portfolio of two sub-transactions.

The first part is a "*loss free loan*" that corresponds to the loan repayment schedule or the lease book value schedule (or book depreciation schedule). This part is loss free in the sense that there will be no shortfall. The duration is until the first stopping time, the earlier of the contractual maturity or default. As we discussed under the Default section, default does not always result in a stop of contractual performance, but since this continuing scenario could be replicated by a subsequent transaction on market terms, we deem that the original transaction has stopped. This loan earns a spread or margin, or lease income (i.e. lease rental, less interest cost, less book depreciation).

Note that only the rental is defined in a lease contract. The interest rate is internal to the lessor as is the rate of book depreciation. Other sources of earnings are possible such as front fees, agency fees, or prepayment fees in loans, or supplemental rent and tax benefits in leases.

The second part is, in the case of a loan, a contingent "*put option*" that the lender has granted the borrower, to put the aircraft to the lender at the stopping time for a strike price equal to the outstanding claim under the "loss free loan". In the case of a lease, the second part is a contingent "*forward purchase*", where the lessor "buys" the aircraft at the stopping time for a price equal to the book value of the aircraft. From a legal point of view, the lessor has been the owner all along, but the "buy" should be seen as re-acquiring the economic ownership of the aircraft. In the real world, there are some added complexities that we could attempt to model to the best of our ability.
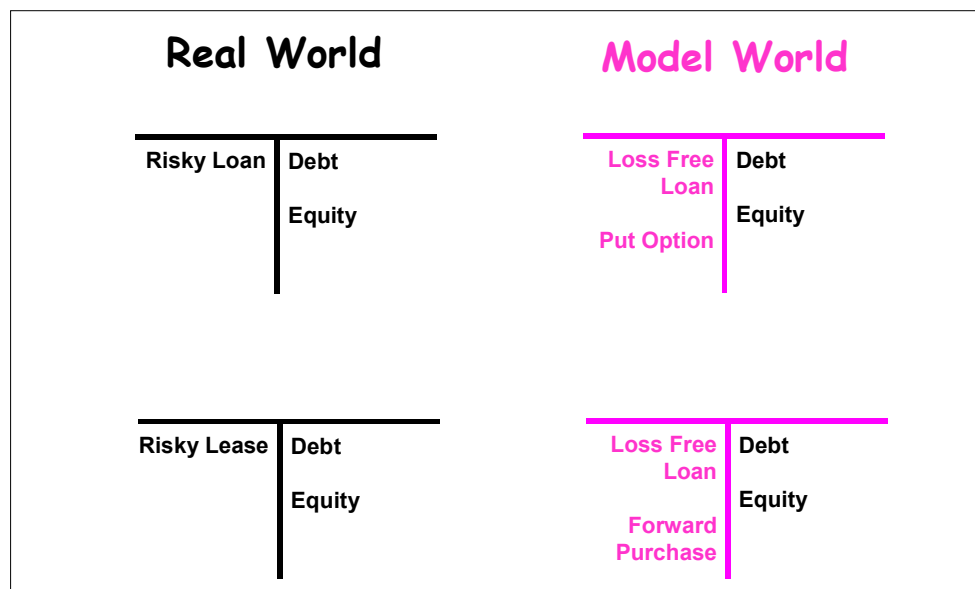

Figure 58: T-account Balance Sheet from the Lender/Lessor vantage point.

The lender will retain an unsecured claim against the borrower for any shortfall, and the lessor will have an unsecured claim for unpaid rent, or at least for the compensating damages in case a subsequent lease would have a lower rental. Recoveries of unsecured claims from airline bankruptcies are anecdotal at best, but often negligible.

A few notes:

1. The pace of book depreciation of a lease does not alter the value of a lease. A faster depreciation will shift value from the "loss-free loan" to the "forward purchase" and vice versa.

2. The outstanding claim and book value above include unpaid interest or rent and any swap breakage cost or gain in fixed interest loans or leases assuming the creditor has matched its funding.

3. Options granted the borrower or lessee to terminate early or extend a transaction always have a negative value for the lender or lessor. There are various ways to calculate this value, but that is outside the scope of this paper.

Subordinated loans can be modeled by designing a portfolio consisting of the loss-free subordinated loan with the put option strike price at the aggregate level of the subordinated and the underlying senior loan, less a put option with a strike price at the level of the underlying senior loan (here, the subordinated lender is the holder of the put option). Leveraged leases can be modeled by designing a portfolio consisting of a loss-free loan, the forward purchase, less a put option with a strike price equal to the claim of the prior ranking claim under the senior underlying loan.

Other types of hybrid transactions, loans with non-recourse balloons, with upside sharing arrangements, or leases with fixed price purchase options, can all be modeled by building a portfolio of sub-transactions. Obviously, the obligor and the aircraft have to be unique in this "portfolio".

The pay-off functions of the "put option" (senior loan) or the "forward purchase" (operating lease) or a subordinated loan can be illustrated in a pay-off diagram as per below. In fact, any hybrid deal can be modeled as long as the pay-off function is monotonously increasing.
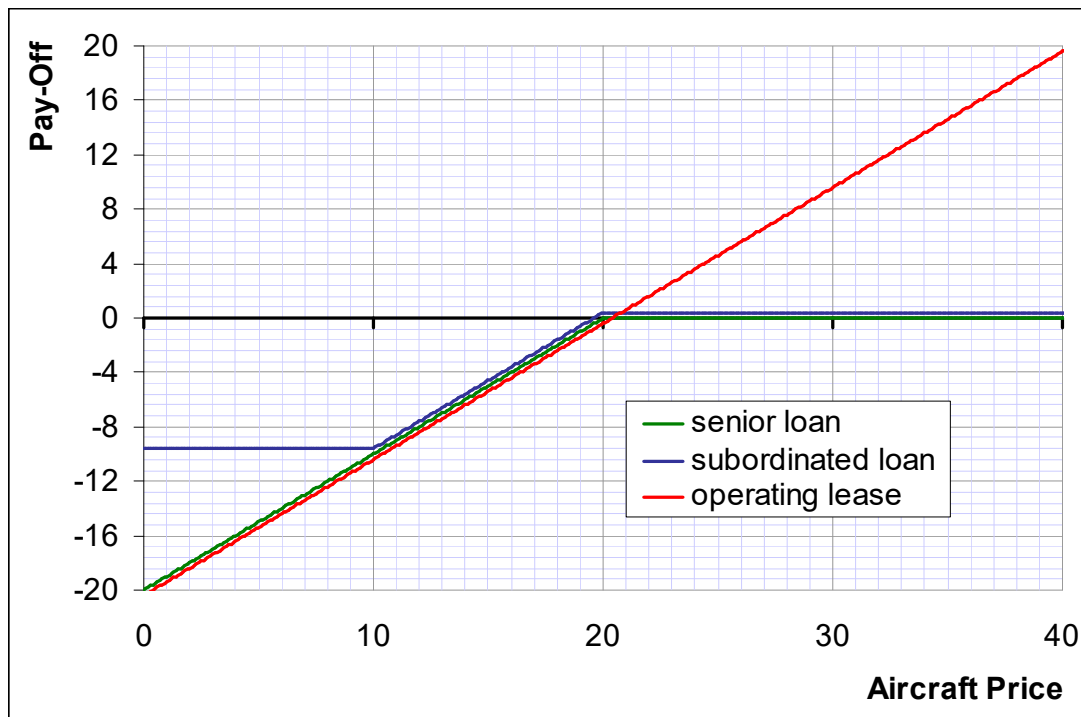
Figure 59: Pay-off as a function of aircraft price for three types of transactions. The creditor's claim is 20 in each case.

Approaches to Modeling Value

The valuation of the "loss-free loan" is simply the net present value of the net payment stream discounted at the risk-free rate calculated up to the stopping time.

The net present value of the "put-option" or the "forward purchase" will be a distribution and will require the calculation of the default probability, the aircraft value path and the interest rate evolution. This can be done in four basic ways:

1) We can use *partial differential equations* that can be solved analytically or numerically;

2) we can use a *measure theoretic approach* to calculate the expectation of the pay-offs;

3) we can build a *Monte Carlo simulation* model to simulate the stochastic variables, and calculate the pay-offs numerically; or

4) we could build an *analytic model*.

The first approach may work for simpler derivatives but is hardly feasible for valuing a transaction as complex as a defaultable equipment backed loan or lease.

The second approach is described in "Asset financing with credit risk" by Vadim Linetsky and Steven Golbeck in the "Journal of Banking & Finance 37 (2013).

Under the first and second approaches, one would apply the tools of financial mathematics that have evolved over the last 35 years. The breakthrough came in 1973 when Fisher Black and Myron Scholes found a closed form to uniquely solve the value of a European call option for a stock under an assumption of constant risk-free rates, constant stock volatility, no transaction costs and continuous trading. Their insight was that the call option could be hedged by a short position in the underlying stock and that this portfolio would be risk-free and would therefore

yield the risk-free rate of return. These approaches applied to aircraft loans and leases would not yield the NPV distributions, but could yield a risk-neutral value that would have to be adjusted for un-hedgeable risk (see page 82-83).

The third approach, called **Monte Carlo** simulation, is the most direct approach and widely used in more or less sophisticated applications. Default outcomes, aircraft value paths and yield curve evolutions are simulated in multiple scenarios, and the pay-offs are calculated directly in each as scenario NPVs. The cycle that alters the expected paths and causes the correlation between default risk, aircraft value, and interest rates is also simulated. The simulated scenarios for each of the stochastic parameters will build up to the assumed distributions. And the scenario NPVs will build up to a resulting NPV distribution.
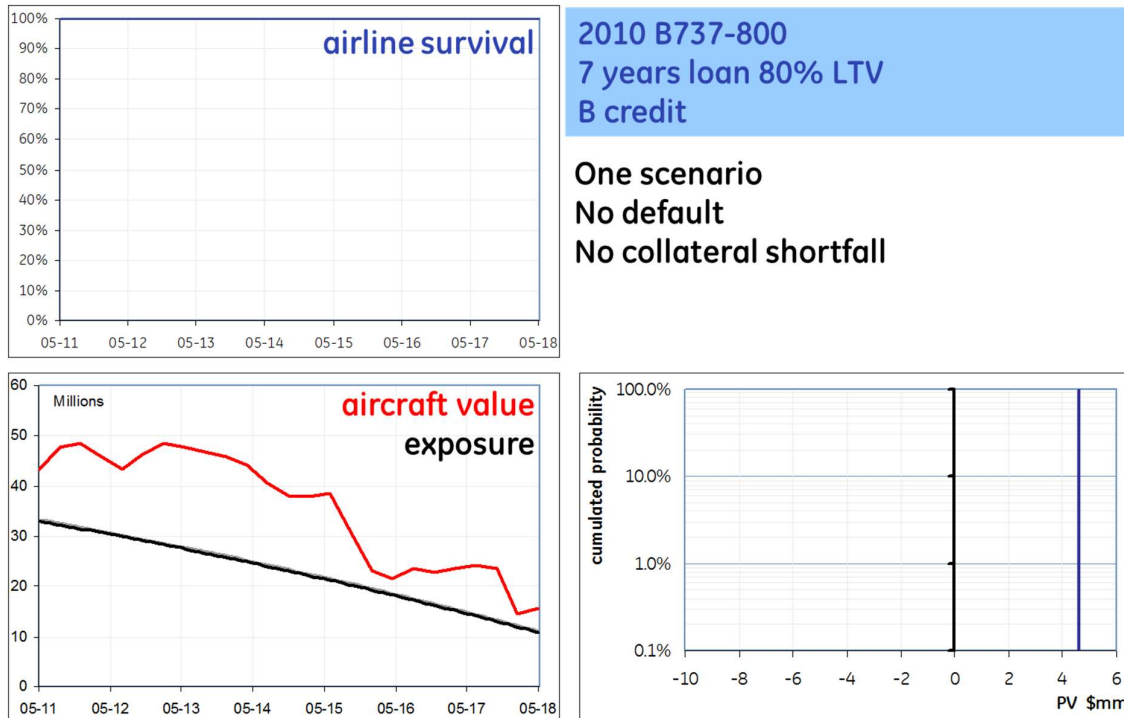
Figure 60: This shows simulated airline performance, the simulated aircraft value path and the resulting NPV number in one simulation of a hypothetical loan. No default occurred in this scenario, and the NPV is the maximum possible: Loan margin for the full tenor and no collateral shortfall.
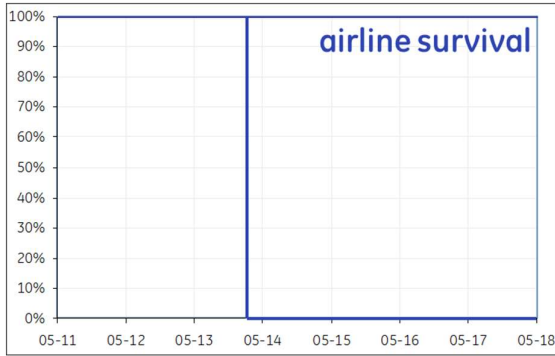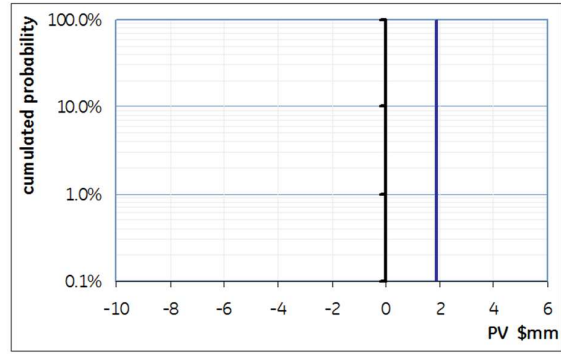
Figure 61: This shows simulated airline performance, the simulated aircraft value path and the resulting NPV number in another simulation of the hypothetical loan. A default occurred in this scenario, and the NPV is less than in the previous scenario, because the margin was not earned for the full tenor. Luckily, the aircraft value covered the exposure at default, so there was no recovery shortfall.
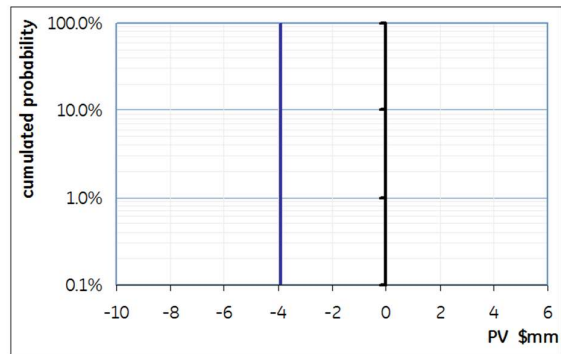
Figure 62: In this scenario, a default occurred, and the NPV is less than in the previous scenario, because the margin was not earned for the full tenor, and the aircraft value did not fully cover the exposure at default. In fact, we suffered a loss.
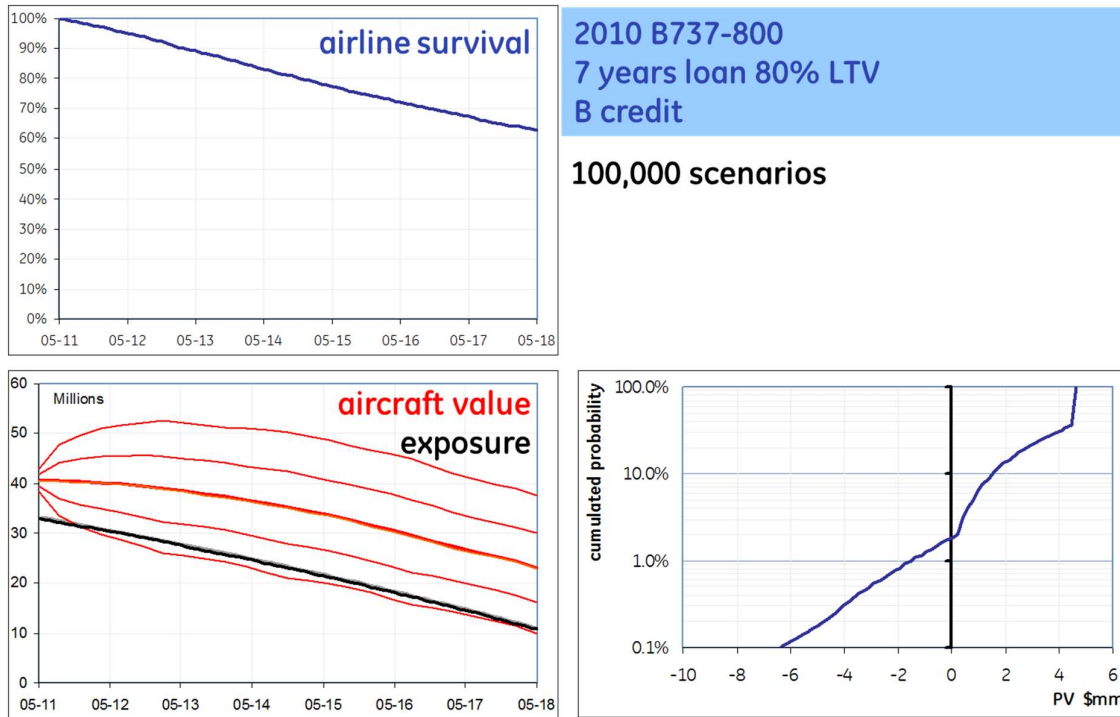
Figure 63:After we have simulated 100,000 scenarios, the airline survival curve and the aircraft value path take on the probability distributions that we have assigned, and the NPV now becomes a probability distribution.

This approach is also very straightforward when we try to analyze a portfolio of several transactions, transactions with multiple aircraft or obligors, or cross-collateralized or cross-subordinated transactions. The drawback is that the computation can be time-consuming.

The fourth approach (building an analytic model) is faster. We take the pay-off function as per Figure 64 (in this case a senior loan) at a particular payment date. We then take the assumed aircraft value distribution at that same payment date as per Figure 65. Note that in this example, the distribution is a Normal distribution, but other distributions may be more representative of the modeling problem (such as an empirical distribution) or behavior of assets (such as the Lognormal distribution). We then take the cumulative probability distribution for every possible aircraft value and plot against the pay-off. This gives us the distribution as per Figure 66.
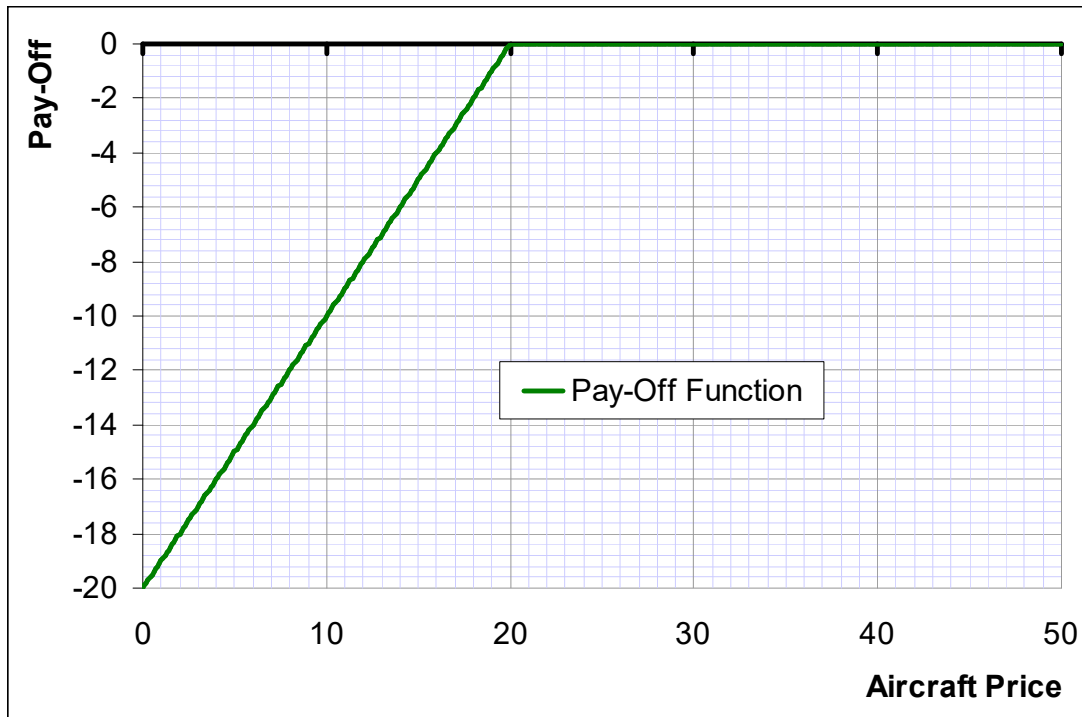
Figure 64: Pay-off function at a specific forward payment date for a senior loan with a claim of 20.
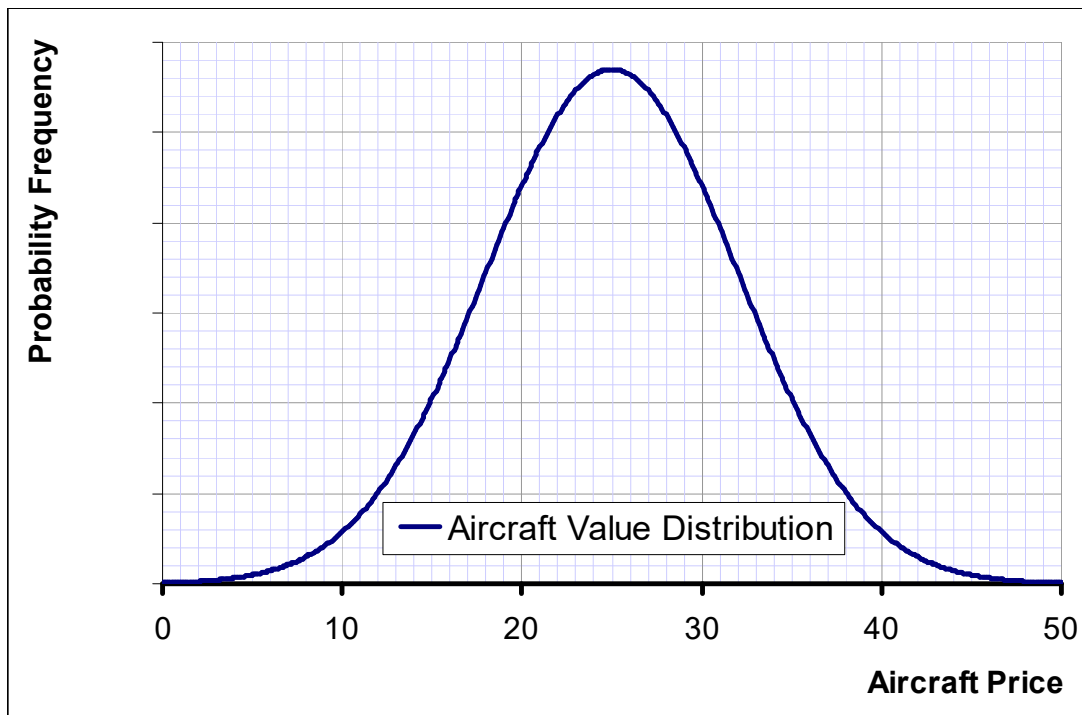
Figure 12 Probability distribution for the value of an aircraft at a specific forward payment date. The expected value is 25 with standard deviation 7. Note that the distribution of the aircraft value in this example is a Normal distribution. When simulating aircraft values paths with a random walk with normally distributed percentage movements, the resulting aircraft value distribution will be Log-normally distributed.

The essence of this analytic approach is the transformation of random variables. There are two random variables involved: the unknown aircraft value and the unknown contract value. Each is represented by a distribution function. The objective is to derive the distribution of the contract pay-off from the pay-off function and the distribution of aircraft values for a given

payment date. This is repeated for each payment date over the tenor of the transaction. See Figure below.
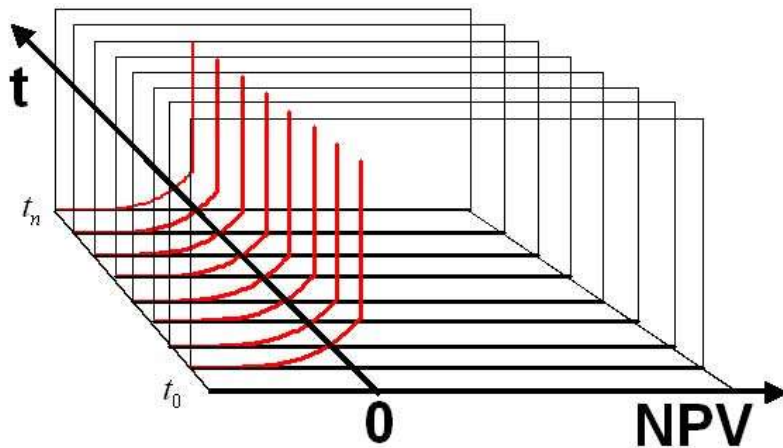
Figure 66: The contractual pay-off distributions for each payment date over the tenor of the transaction are shown in the NPV-Cumulative Probability – time space.

Let us now look at the probability of default at any given payment date. The cumulative probability in this example is shown in Figure . We should model the default probability so that contractual maturity is deemed a default if there has not been a default prior. This way, the "put-option" or the "forward purchase" will be taken into account in cases where we have unamortized exposures without recourse to the obligor (residual, or pure asset value risk). However, if we assign costs for repossession and downtime, those costs should not apply here.
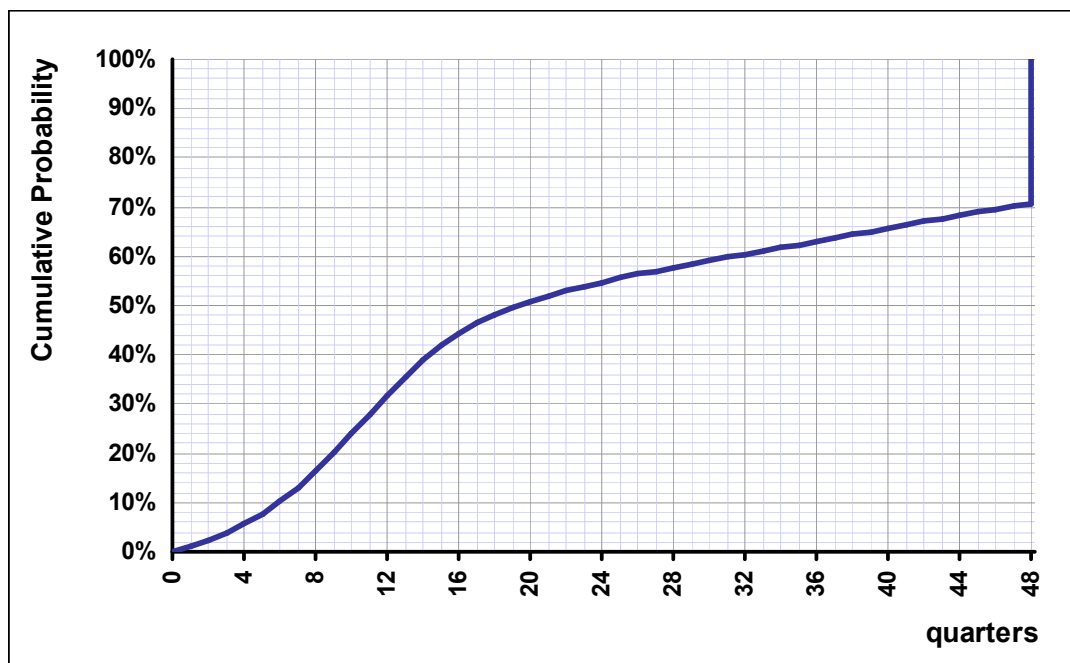
Figure 67: An example of the cumulative default probability for the obligor in this example.

We now take the pay-off distributions for each payment date and discount every value to present value and multiply by the probability of default for that payment period, and aggregate

the results. That will give us the present value of the "put-option" over the full tenor of the loan as illustrated in Figure .
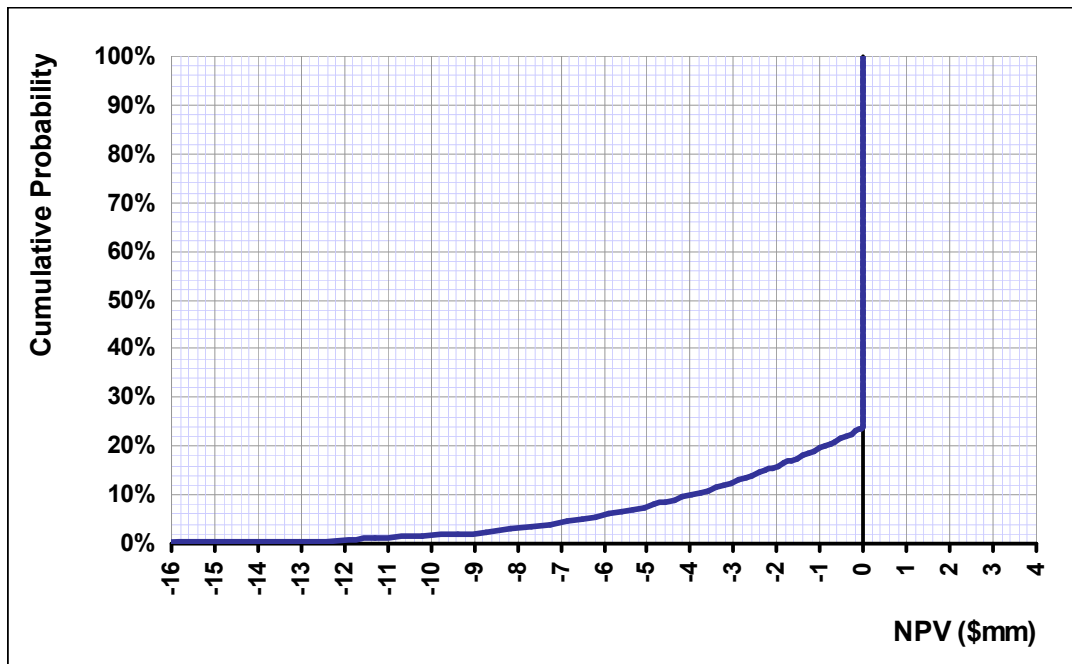
Figure 68: Cumulative distribution for the NPV of the "put option" over the full tenor of the transaction.

Next, we establish the NPV distribution for the "loss-free loan". Remember, we cannot make any losses, but since there may be a default along the tenor, we may make less than the full potential fee and margin income, or in case of a lease, the rental income. The distribution in our example is shown in Figure below.
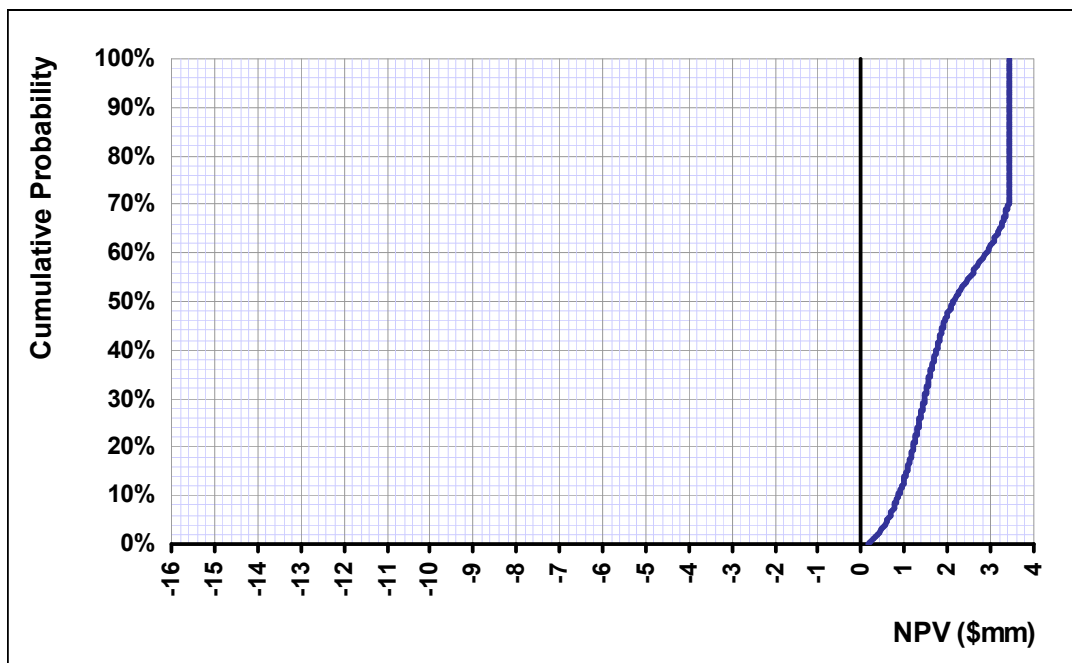
Figure 69: Cumulative probability distribution for the NPV of the "loss-free loan". Note, that only a default will limit the NPV by limiting the tenor below the contractual duration.

When we combine the outcome for the "put-option" and the "loss-free loan", we get the NPV expectation as a distribution shown in Figure 70.
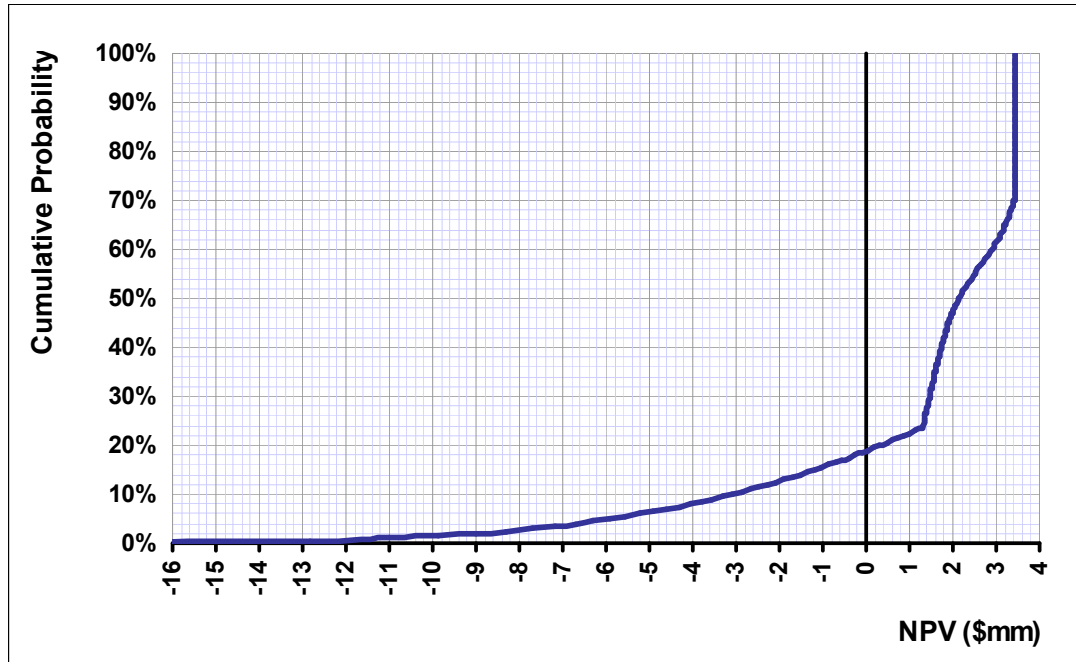
Figure 70: The cumulative probability of the NPV for the combined "put-option" and "loss-free loan".

The drawback of this method is that the cycle that is assumed to cause the correlation among the default risk, the aircraft value, and the interest rates, must be static, and since we cannot possibly forecast the cycle with any meaningful degree of precision beyond a few years, has to be dampened. This resulting lack of correlation effect in the out-years will have to be compensated by arbitrarily wider confidence bands in the parameters.

The process that we have described here is in principle the same as for the Monte Carlo simulation in the previous section, but the resulting distributions in the Monte Carlo simulation can be generated while accounting for full blown cycles in each scenario. On average, the cycle scenarios will dampen out reflecting the uncertainty in the prediction, but the correlation effect will be maintained over the full horizon.

The valuation of the NPV expectation is discussed in the next section.

Risk & Reward

NPV - IRR

The traditional measure to reject or accept a proposed investment/transaction or value is Net Present Value (NPV) or Internal Rate of Return (IRR). The NPV is calculated by discounting all the cash flows with a discount rate that reflects the risk-free interest rate plus a risk premium that somehow reflects the risk in the deal.

$$NPV = -C_0 + \sum_{n=1}^{n=T-1} \frac{C_n}{(1+i)^{t_n}} + \frac{C_T}{(1+i)^{t_T}}$$

where

C_0 is the initial investment,

C_n is the cash flow at time t = n,

i is the discount rate, and

C_T is the balloon for a loan and the residual value at maturity at time T.

If NPV>0, we would accept the deal. The IRR is the discount rate that would result in NPV=0. If the IRR exceeds the risk-free rate plus the risk premium, we would accept.

When IRR is calculated to a zero NPV that has been discounted with the risk-free rate, it is sometimes called Return On Investment (ROI). If ROI is divided by the solvency implied by the capital allocated or leverage applied, we get Return On Equity (ROE). ROI and ROE are classical metrics. The problem is that they are not risk-adjusted and any relationship between ROI/ROE and riskiness is made intuitively or through various grids or tollgates.

In reality, the cash flows are uncertain and the NPV will instead be a probability distribution rather than a fixed number. We can then use the risk-free rate to discount the cash flows from all scenarios and deal with the risk in the shape of the distribution. The IRR distribution is obtained by finding the discount rate that would bring the NPV to zero in each scenario. In the following, we will focus solely on NPV as IRR can be calculated as a spread measure over the outstanding balance we carry over time in an investment.

From our models for forecasting airline default, aircraft value, and interest rates, we can build the various cash flow scenarios that will determine the NPV distribution.

Risk Measures

Once we have the NPV distribution, we have to ask what risk is and how we can measure it? The Concise Oxford English Dictionary defines risk as "hazard, a chance of bad consequences, loss or exposure to mischance". For financial risk, there is no single definition that is entirely satisfactory in all contexts.

Traditionally, financial risk has been expressed as the **standard deviation** of returns. The drawback with this measure is that outcomes above the expectation are counted just as outcomes below the expectation, and those can hardly be thought of as a bad consequence. To address this problem, one could use the **semi-deviation** – calculated just as the standard deviation but only for outcomes on the *downside* of the mean, NPV < Expected NPV. Standard deviation may work if the outcomes are symmetrically distributed around the mean, but this is

not always the case. Especially loans are very asymmetric with a NPV capped at the PV of contractual fees and fixed margin.

Another **downside risk measure** is the conditional mean of all outcomes with NPV<0, multiplied by the probability that a loss will occur at all. We call this measure the **Weighted Average Downside Risk** (WADR). Its mathematical form is derived below with reference to several other, pivotal concepts related to risk-based valuation of a contract.

Five probabilistic concepts and statistics can be identified within this modeling framework:

1. **ECV** – the expected contract value, an unconditional estimate of the contract's value
2. **CECV** – the conditional expected contract value
3. **ADR** – the average downside risk, a special case of the CECV
4. **P(L)** – the probability of a loss (or downside risk)
5. **WADR** – weighted average downside risk

Expected Contract Value (ECV)

This is the unconditional expected value of the contract and could be positive or negative.

$$ECV = E[V] = \int_D v \cdot f(v) dv$$

Where v is contract value and $f(v)$ is the empirical probability density function (epdf) of contract value, considering all component payoff functions. D denotes the domain of integration.

The picture below illustrates the ECV. The shape depicts a typical aircraft backed loan portfolio NPV distribution. (Note that here, we are showing the probability frequency distribution as opposed to the cumulative probability distribution.) The pink portion represents the outcomes below zero when we have losses. The white portion represents the positive outcomes. If we cut out the shape in card board and attempted to balance the shape on a ruler parallel to the

vertical axis (the grey line), it would balance exactly at the point representing the Expected Contract Value (or expected NPV). This is the mean of all outcomes.
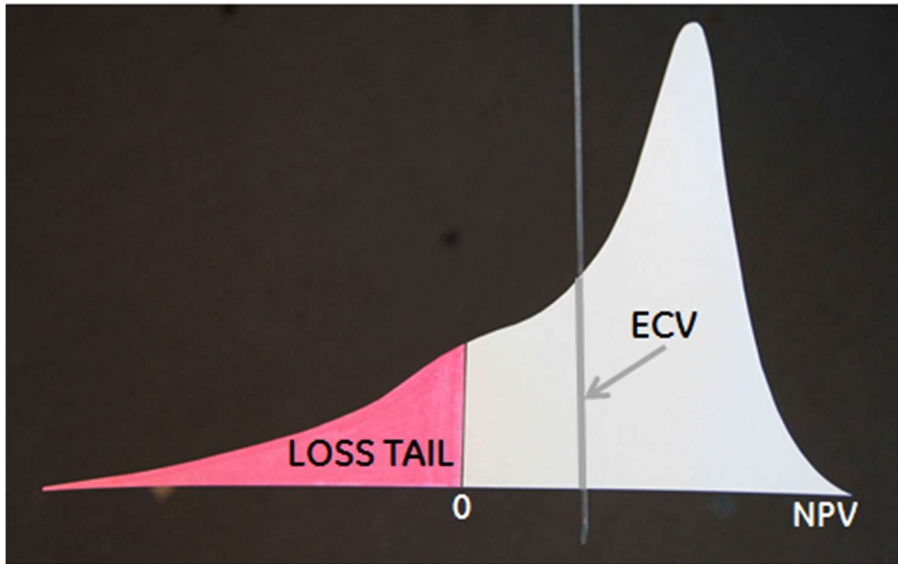
Conditional Expected Contract Value (CECV)

It is possible to take the mean of any subset or 'region' of the contract's value when it is expressed as a random variable. Mathematically, the CECV has a the following structure:

$$CECV = E[V \mid V < v^*] = \frac{\int\limits_{-\infty}^{v^*} v \cdot f(v)dv}{\int\limits_{-\infty}^{v^*} f(v)dv}$$

When v^* equals zero, we are computing a specific case of the CECV which we discuss below. This is the basis behind two risk measures for assessing loans and leases, the Average Downside Risk (ADR) and Weighted Average Downside Risk (WADR). The ADR is a special case of the CECV.

Average Downside Risk (ADR)

The ADR is the conditional mean value of the contract given that a loss has occurred and is expressed as follows:

$$ADR = \left| E[V \mid V < 0] \right| = \left| \frac{\int\limits_{-\infty}^{0} v \cdot f(v)dv}{\int\limits_{-\infty}^{0} f(v)dv} \right|$$

Using the same example as for the expected NPV above, the picture below illustrates the tail cut out from the cardboard NPV distribution. The line where this shape balances on the ruler is the ADR, or **Conditional Expected Loss**.
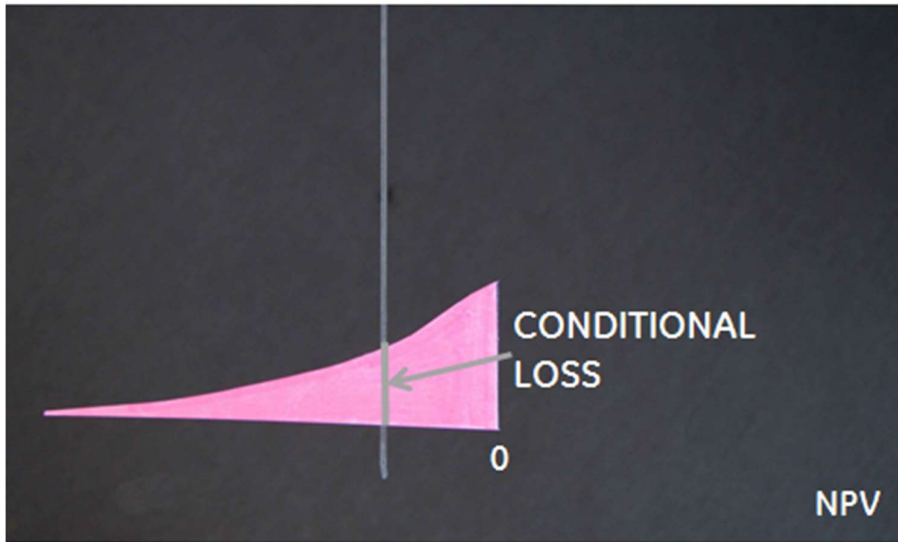
Probability of a Loss (P(L))

This is probability of the realized contract value (at time t) being less than zero and is equal to the empirical cumulative distribution function evaluated at $v = 0$.

$$\Pr(V < 0) = F(0) = \int_{-\infty}^{0} f(v)dv$$

Weighted Average Downside Risk (WADR)

This is the mean loss, given that loss has occurred, multiplied by the probability of a loss occurring at all.

$$WADR = ADR \times Pr(L)$$

$$WADR = \left| E[V \mid V < 0] \right| \times \Pr(V < 0) = \left| \frac{\int_{-\infty}^{0} v \cdot f(v)dv}{\int_{-\infty}^{0} f(v)dv} \right| \times \int_{-\infty}^{0} f(v)dv = \left| \int_{-\infty}^{0} v \cdot f(v)dv \right|$$

This may be a bit less intuitive. To continue the analogy with our NPV distribution cut out in card board, we would now bend the shape along a line parallel to the vertical axis where losses are to the left and gains to the right (NPV=0) in a 90 degree angle. (This means all the mass from the positive outcomes falls on the NPV=0 axis.) The WARD, or the **Unconditional**

77

Expected Loss, is the point where the bent shape balances on the edge of the ruler as seen in the picture below.
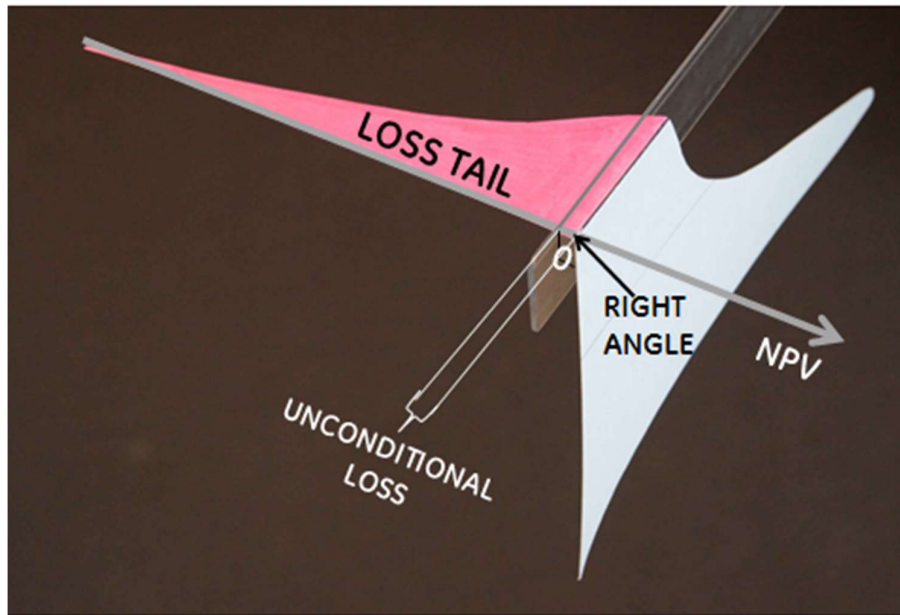
In technical terms, the WADR is a special case of a first lower partial moment (FLPM)[1] where the reference point *(q)* is zero. The general form (using *V* as the key random variable) for k^{th} order partial moments is as follows:

$$LPM(q;k) = \int_{-\infty}^{q} (q-v)^{k} \, dF_{V}(v) \, dx$$

The WADR is a risk measure that could be used to estimate a reserve to be held against a specific contract in a portfolio of contracts. It is analogous to the expected loss (EL) in the Basel II framework:

Expected Loss = PD × LGD × EAD

Note that the Basle II framework looks at the loss distribution alone, not the full NPV distribution. If all gains were ignored, ADR and WADR would be the same, and PD and P(L) would basically be the same.

The expected loss is often a component of performance measures for pricing a loan, where it is subtracted from net income before being divided by economic capital. Risk-Adjusted Return on Capital (or RARoC) is an example of such a measure:

RARoC = (Net Income – Expected Loss) / Capital

To meet a target RARoC (i.e. the hurdle rate)—and given an Expected Loss computed from the attributes of the transaction—the margin and fees must be set to cover the risk of the transaction.

---

[1] Danielsson, J., Zigrand, J-P., Jorgensen, B. N., Sarma, M., and de Vries, C. G., (2006). *Consistent Measures of Risk,* Discussion Paper 565, Financial Markets Group, London School of Economics and Political Science.

In the context of Basle II, the **Expected Loss** is almost the same as Weighted Average Downside Risk over a one-year horizon (except that the Expected Loss is not mitigated by the expected income over that one year period). The WADR is a natural measure for loss reserves in an expected loss approach. One can think of WADR as the probability of loss times the magnitude of loss. Reserves of this amount would cover losses *on average across a whole portfolio and over time*. This measure gives a sense of the overall or **average** "riskiness" of a transaction.

Please note that the Expected Loss is not a loss expected. Transactions where we expect to make money still have a positive Expected Loss. Expected Loss is merely a measure of the loss we expect to make *if we make a loss*, multiplied by the probability that we will make a loss.

Besides making less money than we had hoped for, or losing some money, the one thing that people don't like is to lose a large amount of money—to be ruined. We can think of several different **bad-case** risk measures. One common measure is **Value at Risk** (VaR) which is the loss that will not be exceeded with a given confidence, such as 95% or 99% or 99.99%.

We take V to be a random variable describing the value of a contract, with the associated density and distribution functions as described earlier. There are two types of VaR: *absolute* (loss levels relative to zero) and *relative* VaR (loss levels relative to the mean or expected return). We focus on *absolute* VaR as it the most relevant for the framework under discussion. For a given confidence level α, we can define VaR_α in the following way:

$$VaR_{\alpha,t} = q = Q_{(1-\alpha)}(V_t)$$
$$= F_{V(t)}^{-1}(1-\alpha)$$

In other words, the absolute VaR is the $(1-\alpha)^{th}$ quantile of the cumulative distribution function for the total contract value at t.

One shortcoming with this measure is that it doesn't say anything about how bad a loss could be when VaR is exceeded. To address this, we can measure the probability weighted NPV below VaR, called Conditional VaR , **CVaR** or Expected Shortfall.

$$CVaR_{\alpha,t} = \left| E[V_t \mid V_t < F_{V(t)}^{-1}(1-\alpha)] \right|$$
$$= \left| E[V_t \mid V_t < q] \right|$$
$$= \left| \frac{\int_{-\infty}^{q} v \cdot f_t(v)dv}{\int_{-\infty}^{q} f_t(v)dv} \right|$$

A problem with these loss measures is that they are calculated at the extreme tail of the NPV distribution, in the region where we often have very little empirical data to validate the shape. As an example, a confidence interval of 99.99% means that we would only expect the loss to be exceeded one time in 10,000 trials. There just aren't enough aircraft loans or leases, or years of history in the jet age to collect a meaningful amount of data points that would validate our view of the shape of the tail in this region. Also, the purpose with these extreme loss measures is to avoid ruin and calculate an adequate amount of capital to shield us. But capital should rather be set aside on a portfolio level and the impact of a single transaction should be based on the marginal capital that would need to be added to the whole portfolio. This requires us to

model the NPV distribution of the full portfolio and accounting for interdependencies between different aircraft and airlines. More about that on pages 94 -108.

It would be unwise to rely blindly on VaR or any other bad-case risk measure as a basis for allocating capital or determining an appropriate leverage. At best, we can maybe build a model that will be directionally sound, in other words, a model that will help us allocate capital in relative terms, across transactions, portfolios, and over time. And that is not too bad!

There is also another aspect to capital allocation: We are not setting capital for a transaction, but for a portfolio or a business segment. Our need for capital will not only be determined by the outcome of isolated transactions, but on the performance of **a continuous business activity**. It may well be that a period of frequent defaults with corresponding aircraft value shortfalls will be followed by a period of very attractive trading conditions that will generate excess returns.

Consider the example of a defaulting lessee returning an aircraft with a market value well below the book value in a down-cycle, thus generating an economic loss for the lessor. The lessor may lease this aircraft on a new lease and the market value may pick up over the ensuing lease period. The economic loss is now off-set by an economic gain. The lessor's ability to raise financing to fund this ongoing business process is as much a determinant of survivability as the economic loss in a cluster of individual transactions. Obviously, the two are related, but liquidity in the capital and banking markets is crucial.
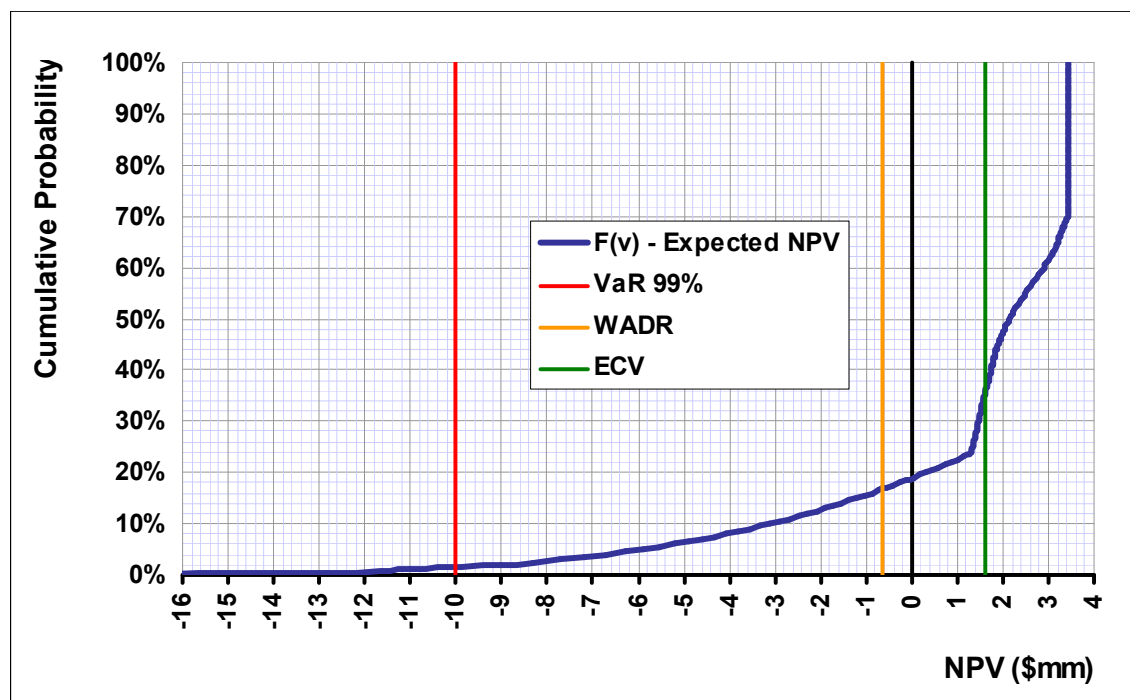
Figure 71: The blue line represents the cumulative probability of an outcome below the corresponding NPV. The green vertical line shows the expected, or mean, NPV. The orange vertical line is the sum of all losses multiplied by their probability of occurring, the Weighted Average Downside Risk. The red vertical line is the loss that will not be exceeded with 98.5% confidence – 1.5% on the vertical probability axis, Value at Risk (98.5%).

The academic literature suggests that a **coherent** risk measure should meet four criteria[2]:

---

[2] Artzner, P., Delbaen, F., Eber, J.M., Heath, D. (1999). Coherent measures of risk. Math. Finance 9(3), 203-228

1) *translation invariance* (holding a cash amount should reduce the risk by exactly that amount);

2) *sub-additivity* (the risk in a portfolio of two transactions should not exceed the sum of the risk of the two transactions individually);

3) *positive homogeneity* (when the size of the position doubles, the risk doubles too), and;

4) *monotonicity* (if one portfolio always produces a worse outcome than another, its risk measure should be greater).

Neither WADR, nor VaR meet all criteria all the time. WADR is not translation invariant unless positive earnings are disregarded. VaR is not always sub-additive. Yet, they are intuitively appealing. The WADR measure is appealing because it represents the right loss reserve based on expected loss. VaR is appealing because it sets a capital amount that would protect against ruin with a given confidence, just as a rating represents a probability of default.

How Do We Price for Risk?

We have looked at the downside of the NPV distribution to measure risk. This is helpful to set loss reserves and allocate capital. We can look at the full NPV distribution to set rational selection criteria for new investments and to value individual transactions.

The mean of the distribution is the Expected Present Value. Is that the value of the transaction? Suppose we had an aircraft lease or an aircraft loan that had an expected return of 10%. Would we consider that investment to have equal value to a treasury bond with the same return? Probably not. We know that the aircraft lease or loan has uncertain returns with a risk of making less than 10% as well as more. For a given expected return, we would prefer the investment with less uncertainty. By the same token, for a given risk, we would prefer the investment with the higher return. This holds true except for risk seekers who attach some kind of value to the thrill of uncertainty – or how would one otherwise explain gambling?

But if we look strictly at monetary returns, an investor would seek to be compensated for taking risk over and above the expected loss. However, many investments with uncertain returns such as stock options and other derivatives, trade at a price that basically just compensates for the expected loss. This is *called risk-neutral pricing*. The reason why there is no premium is that the investor can hedge, at least in theory, the derivative so that the portfolio becomes risk free. So for example could a buyer of a stock option go short in a portion of the underlying stock so that a small change in the stock price would be exactly off-set by the change in the value of the option. This is the basis for valuing financial derivatives provided they trade in a *complete market*. The completeness means that one can perfectly hedge the position. In more general terms, the capital market theory holds that only risk that cannot be hedged or diversified away will command a premium.

For aircraft loans and leases, the market is incomplete, and the risk-neutral pricing is not a correct basis for valuation. There is no generally accepted theory how to price for risk in an incomplete market.

There are three basic approaches we can take: 1) The actuarial method, 2) the utility indifference pricing method or 3) the financial method.

In **actuarial pricing** (borrowed from the insurance industry), the pricing should cover a) administrative costs, b) expected losses (average losses), and c) a risk premium. The risk premium is calculated as a hurdle rate on the economic capital required as a buffer for unexpected losses. The economic capital could be calculated as any of the bad-case risk measures above. This method is an *absolute* pricing method and may be workable for loans and leases that are not traded and for which there are no market prices. The traditional bank methods come under different names such as RARORAC (Risk Adjusted Return On Risk Adjusted Capital) or RAROC or RORAC. However, the hurdle rate is arbitrary as is the method for setting the level of economic capital allocated to the transaction.

In **utility indifference pricing**, we look at an investor's risk aversion or risk tolerance. What is the amount of cash in hand that would make an investor indifferent to a risky investment? This is an *internal* pricing method that depends on the investor's utility function. The utility function is assumed to be increasing to reflect that investors prefer more wealth to less, and concave because investors are risk averse. The drawback with this pricing method is that different investors have different utility functions, so the price is not unique.

In **financial pricing**, we define a model for the risk premium and calibrate the parameters to observed market prices. This is a *relative* pricing methodology and requires that we have observable market prices. Since no two deals are identical, the market price that we observe for one is not directly applicable to another. It will be unclear if it is measures around airline default, aircraft value and volatility, interest rate movements, and term structure on one hand, or our assumptions about the risk premium on the other, that is being calibrated to the market. Also, very few aircraft loans and leases are traded, at least publicly, so the market data is scant. The EETC and ABS markets will be helpful, but having a broad and deep presence in the market will be even more helpful. This gives large lenders and lessors an edge of superior information. The more different aircraft types and airline credits we are involved with, the more transactions we do, the crisper the risk premium picture will become.

One variant of financial pricing would be to price to risk-neutrality, but adjust the probabilities for aircraft values, default risk, and interest rates from the real world to the risk neutral world. In other words, we would stress the simulated aircraft values paths, obligor default probabilities and the yield-curve evolution to a point where the risk neutral pricing would become consistent with market prices. There are two advantages with this approach; 1) the calibration against market prices could be made to traded securities whose values depend directly on aircraft values (although we cannot think of such a security right now), default risk or interest rates, and 2) the kit of mathematical tools and methods that have been developed for complete markets could be used.

How do we then define a model for the risk premium? In investment management, performance is sometimes measured with the Sharpe Ratio. This is the ratio between the return over the risk-free return divided by the standard deviation of returns. A strong investment performance means high returns with low standard deviations. As we discussed above, given the asymmetric shape of the return distribution for most loans and some leases, WADR may be a better risk measure. We can easily translate a Sharp Ratio to a performance measure based on WADR as a risk measure. One such measure is WADR/Expected NPV, that we will call the Risk Reward Ratio (RRR). This measure is discontinuous when the Expected PV becomes zero, something that makes it unsuitable for transactions with low or negative expected returns.

Mathematically, the RRR can be expressed as

$$RRR_1 = \frac{WADR}{ECV} = \frac{\left| \int_{-\infty}^{0} v \cdot f(v)dv \right|}{\int_{-\infty}^{\infty} v \cdot f(v)dv}$$

An alternative is to use WADR divided by the Weighted Average Upside Opportunity (WAUO), calculated just as WADR, but just taking the probability weighted portion of the NPV distribution when NPV>0. This is specified as follows:

$$RRR_2 = \frac{WADR}{WAUO} = \frac{\left| \int_{-\infty}^{0} v \cdot f(v)dv \right|}{\int_{0}^{\infty} v \cdot f(v)dv}$$

This "downside/upside" ratio goes from 0 for a very good transaction to infinity for a poor one.

If we take the inverse of the latter ratio, it can also be written as the Omega for L=0.

$$\Omega(L) = \frac{\int_{L}^{\infty} [1 - F(v)]dv}{\int_{-\infty}^{L} F(v)dv}$$

Where $F(v)$ is the cumulative distribution function of the value of the contract; alternatively, the distribution of returns. This measure goes from 0 for a poor transaction to infinity for a good one. Omega is an increasingly popular investment performance measure that can be seen as a ratio between the value of a call option and a put option with a strike price of L[3].

The key is to grasp what the ratio should be. Clearly the benchmark ratio changes over time depending on the appetite for risk in the market. If we could pin-point the market ratio, we could move our NPV distribution along the NPV axis to a point where the ratio equals the market (or benchmark) ratio. The more active we are in the market, the better can we develop a view of what the current market RRR is in our investment space. We may also look at market RRRs over long periods, and be able to identify when the market seems to overheat – "irrational exuberance" maybe? This translation would represent the mark-to-market value (M-t-M) of the transaction and indicate a premium or a discount we could expect when we sell or buy a loan or a lease.

---

[3] Keating, C., and Shadwick, W. F., (2002). A Universal Performance Measure. *The Finance Development Centre (London)*. Web Ref: http://www.actuaries.org.uk/__data/assets/pdf_file/0004/19480/keating.pdf.
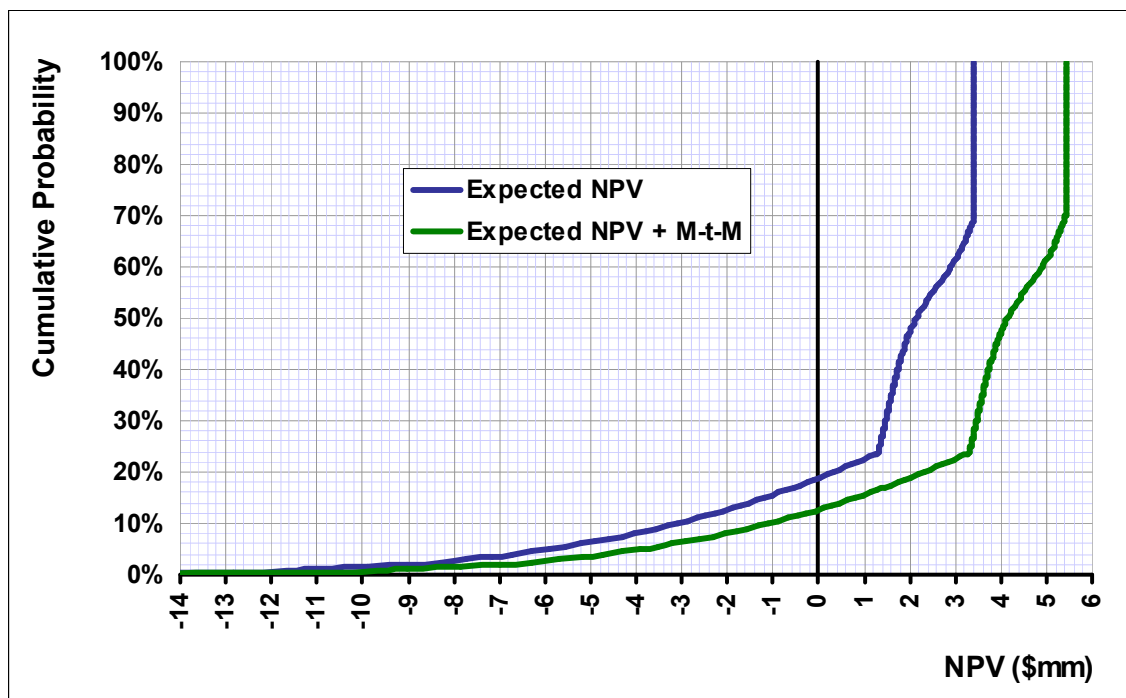
Figure 72: Cumulative probability distribution for the NPV for a transaction with a RRR above the benchmark. By adding $2MM cash to the position, the distribution now shifts to the green line which reaches the benchmark RRR.

As the market is rather illiquid and private, it would maybe be more appropriate to call the M-t-M the "Mark-to-Model" value. Although M-t-M is not a risk measure, it meets all the four coherence criteria as per above.

Another way to illustrate the M-t-M is to show where our transaction falls in a WADR – NPV chart. The blue line in Figure represents the current market RRR. The green dot shows where our transaction falls. The M-t-M (in this case positive) is the vertical distance on the NPV axis above the intercept with the RRR market line. The red dot represents the transaction with a negative cash position. As the cash position changes, the transaction moves along the red line. This line is curved and is unique for each transaction depending on the properties of the NPV-Probability distribution.
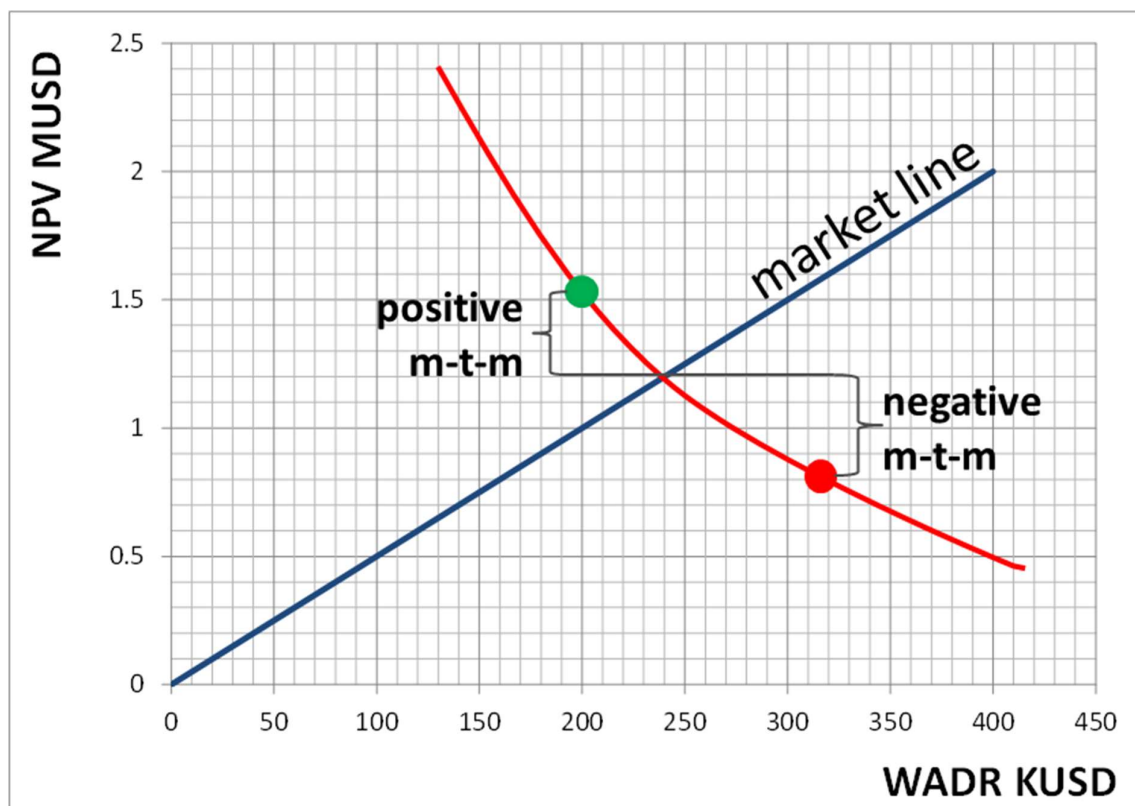
Figure 73: A loan or lease positioned in the NPV – WADR chart with a market RRR line. The Mark-to-Model value is the vertical distance from the RRR line.

It is not clear that Risk-Reward Ratios or Sharpe Ratios are consistent across different asset classes and that there would indeed be a "market" RRR at any given point in time. Reducing a complex two-dimensional picture of the NPV distribution to a single ratio is less than ideal; however, it is a necessary step towards ranking and selecting amongst alternative investment opportunities. At best, we should seek a coherent and consistent pricing and valuation methodology for our asset class of aircraft loans and leases.

In addition to risk, investors also seek to be compensated for the lack of liquidity in an investment. Publicly traded securities can be bought and sold in real time and at little cost. Aircraft loans and leases are time consuming to trade and the transaction costs can be substantial. This lack of liquidity will also affect price.

Risk Premium

Once we have determined the M-t-M, it is easy to define the Risk Premium. The Risk Premium is the difference between the Expected Present Value and the M-t-M.

$$RP = ECV - MtM$$

This Risk Premium can be translated to a spread by dividing by the present value of the Balance Volume

$$BV = \sum_{i=1}^{t_M} Outstanding\ Principal_{(i-1)} \times \frac{t_i - t_{i-1}}{365.25} \times e^{-r_i t_i}$$

We have discussed the M-t-M of a transaction. A favorable market value should be a sufficient reason to originate a transaction, but it doesn't automatically follow that we should hold it. The internal value may be different from the market value for several reasons:

1.   the transaction may add excessive obligor or aircraft type concentration to the portfolio;

2.   it may add too much sector concentration to our institution;

3.   our tax benefits from holding the transaction may be inferior to the rest of the market; or

4.   our cost of funds (or cost to manage) may be higher than the rest of the market.

These points may be true of any transaction in our portfolio. As the portfolio evolves—and the circumstances of our institution evolve—the internal value relative to the market value will continuously change. Keeping track of this is the basis for proactive portfolio management. Clearly, the benefits of active syndication and asset trading must be weighted against the transaction costs, and it is crucial to structure loans and leases to reduce those.

The M-t-M is an appealing metric for tracking origination and portfolio management performance, and may provide a better alignment of incentive compensation to shareholder value than some volume or net income based schemes.

Example of a Unit Loan

NPV-Probability Distribution for the Base Loan

Let's now illustrate how the NPV-Probability chart is affected by various properties of a loan. We will use a Unit Loan, a building block with the same term structure in all examples. For simplicity, we assume no transaction costs.
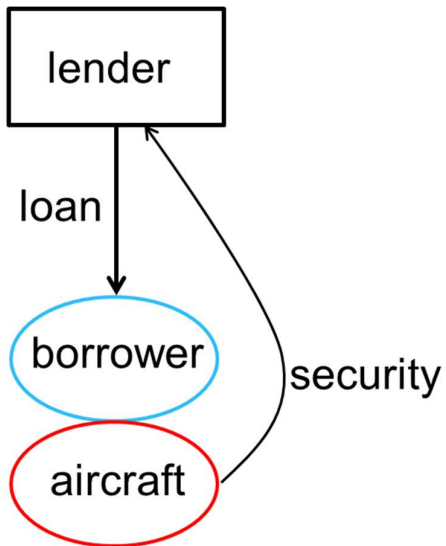
Figure 74: A schematic for the Unit Loan. This is a simple one borrower – one aircraft loan with full recourse to the borrower and a set term structure.

The chart below shows the Net Present Value on the horizontal axis and the cumulative probability on the vertical axis.
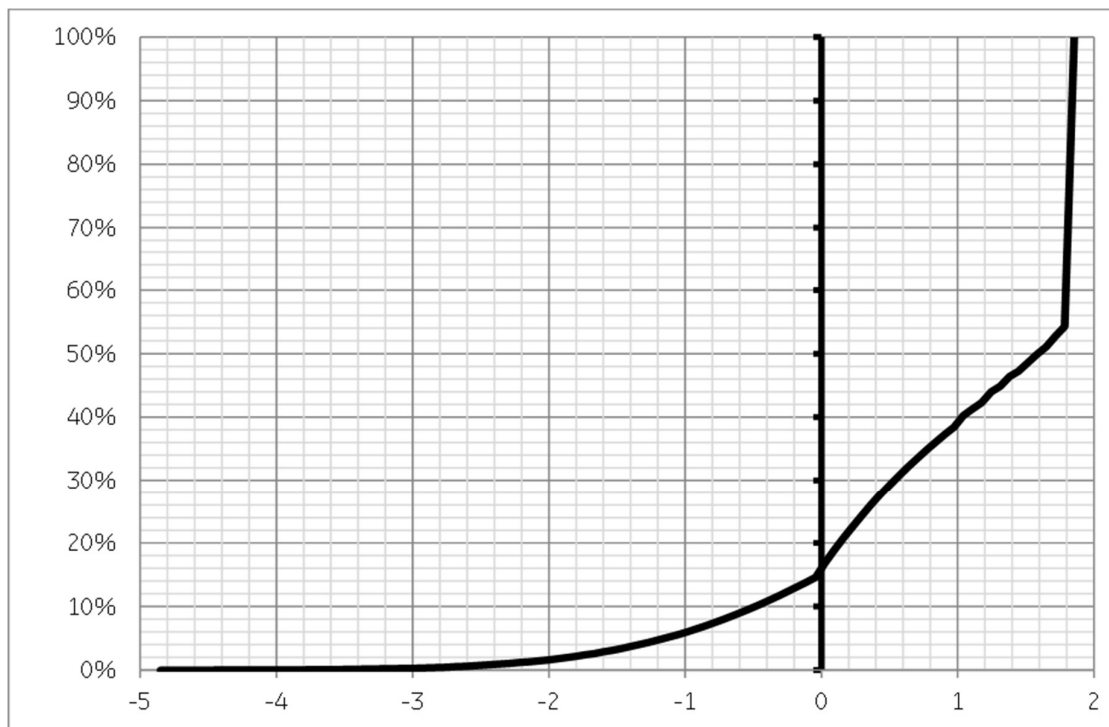

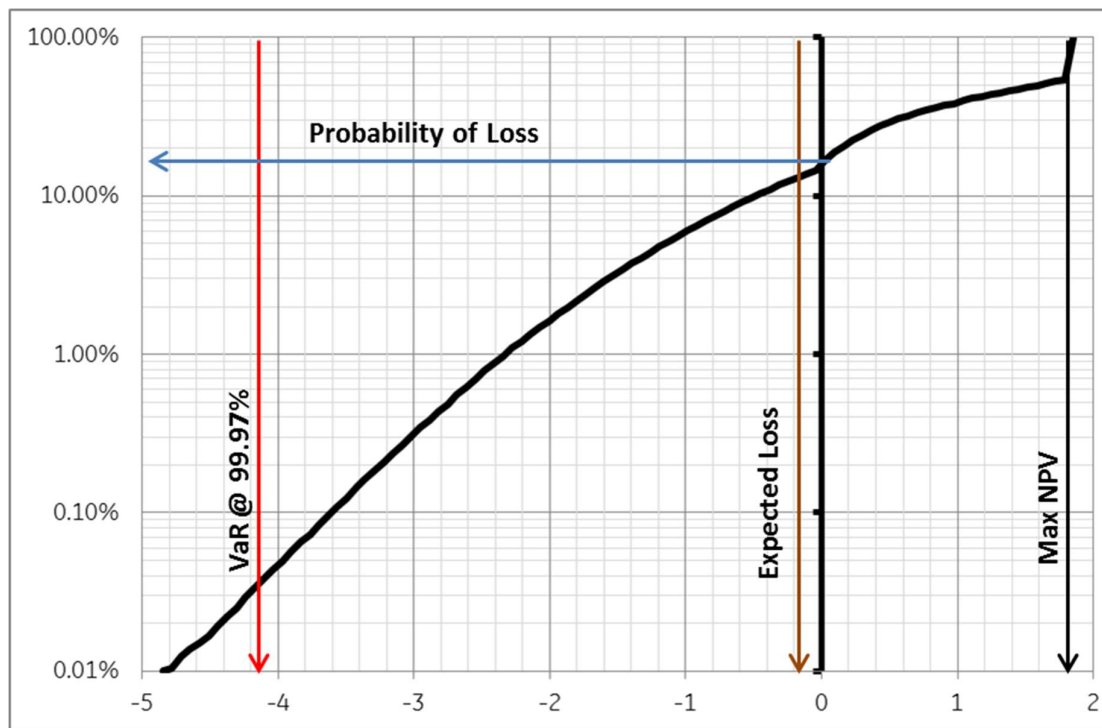*Figure 75: The NPV-Probability distribution for the "Unit Loan".*

Figure 76: The NPV-Probability Distribution for the "Unit Loan" with a logarithmic probability scale. This way, it is easier to see the "tail" of the distribution.

The chart above shows the cumulative probability that the Net Present Value will fall below the value on the horizontal axis. The probability is plotted on a logarithmic scale to show the loss tail better. In this case, with about 43% certainty, the obligor will survive the full term, and we will make the maximum potential Net Present Value (1.85mm). In scenarios where the obligor defaults prior to maturity, the NPV will be less than the maximum, first because margin income will be foregone and secondly because we may have a recovery loss (the aircraft being worth less than the outstanding balance).

The probability of losing money (having a negative NPV) is 16.5% in this example. The Expected Downside Risk (the magnitude of all losses [an NPV below zero] multiplied by the probability that that scenario will occur) is 0.145mm. The Value at Risk at the 99.97% confidence level (only a 0.03% probability that the loss will exceed this value) is -4.15mm.

Impact of Obligor Rating

The next chart shows the effect on the NPV distribution of the counterparty rating of the obligor of the Unit Loan. The base case is a B rating (black line), the red lines show ratings of CCC and CC, while the green lines show BB, BBB, A, AA, and the Immortal obligor as the vertical green line. In the latter case, the outcome is always the maximum potential NPV – no default and no collateral shortfall.

The dotted red line shows the distribution in the case of immediate default, no potential for a positive NPV. This NPV distribution dominates some of the lower rating grades for a small part of the probability spectrum. The reason for this is that although the expected LTV profile for the Unit Loan is declining, the uncertainty of the aircraft value increases over time.

VaR and Expected Losses decrease with improved ratings – however, that may not always be the case: Loans with very low initial LTVs and slow amortization and/or non-recourse balloons may well have an inverse obligor rating to loss relationship.

Please note that the Unit Loan always has the same loan margin – in reality, a loan with a better obligor rating typically commands a lower spread.
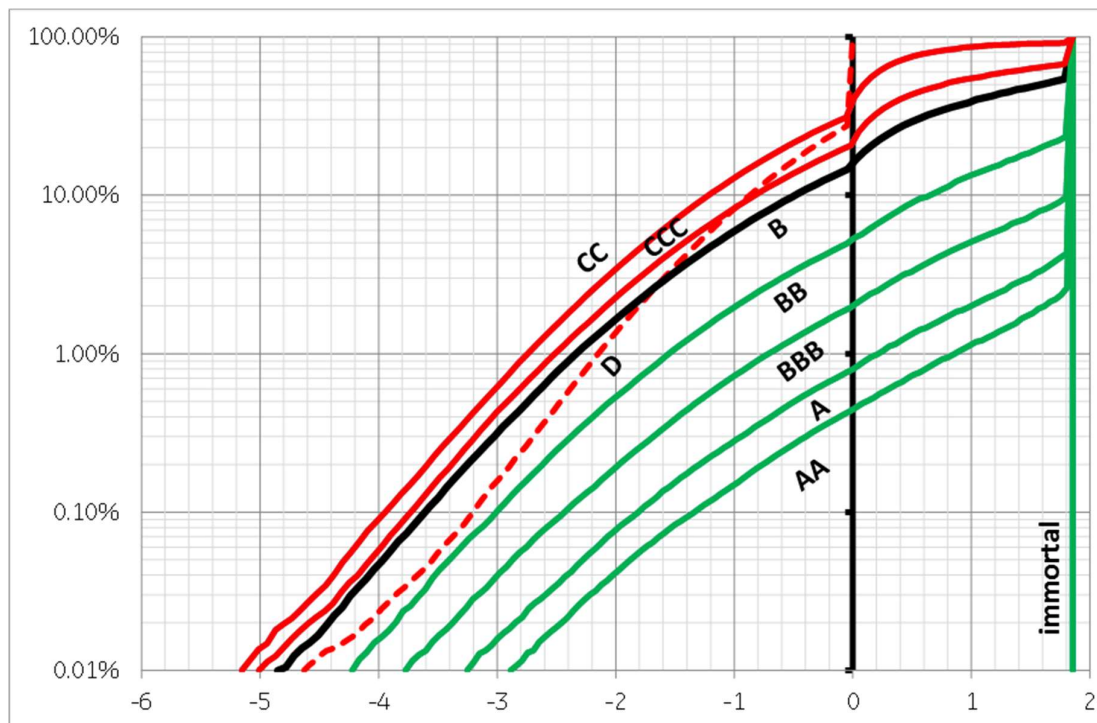
Figure 77: The Unit Loan with the counter-party rating varied.

Impact of the Loan to Value Ratio

The chart below shows the effect on the NPV distribution of the Loan-To-Value Ratio of the Unit Loan. As the loan amount has been reduced pari-passu along the amortization profile to achieve the displayed LTV, the amount has been re-scaled to always represent a 10mm initial loan amount.

The black line represents the Unit Loan with red lines showing 90%, 100%, 110% LTVs, while the green lines represent 70%, 60%, 50%, 40% and 30% respectively. The dotted gray lines show the outer boundaries, infinite LTV (unsecured) and zero LTV (loss free). The zero LTV case obviously shows no loss since the exposure is infinitely covered, but the B rating assigned to the obligor in the Unit Loan means that early termination due to default risk will lead to foregone earnings. The infinite LTV case corresponds to an unsecured exposure with a zero recovery in case of default.

VaR and Expected Loss are mitigated by lower LTVs, with the lower LTV always dominating.

Please note that the Unit Loan always has the same loan margin – in reality, a lower LTV loan typically commands a lower spread.

Figure 78: The Unit Loan with the initial LTV ratio varied.

The Impact of Subordination

The effect of subordination is shown in the chart below. The Black line is the Unit Loan. The blue lines represent junior ranking parts of the Unit Loan, re-scaled to the same amount as the Unit Loan, with different levels of sub-ordination.

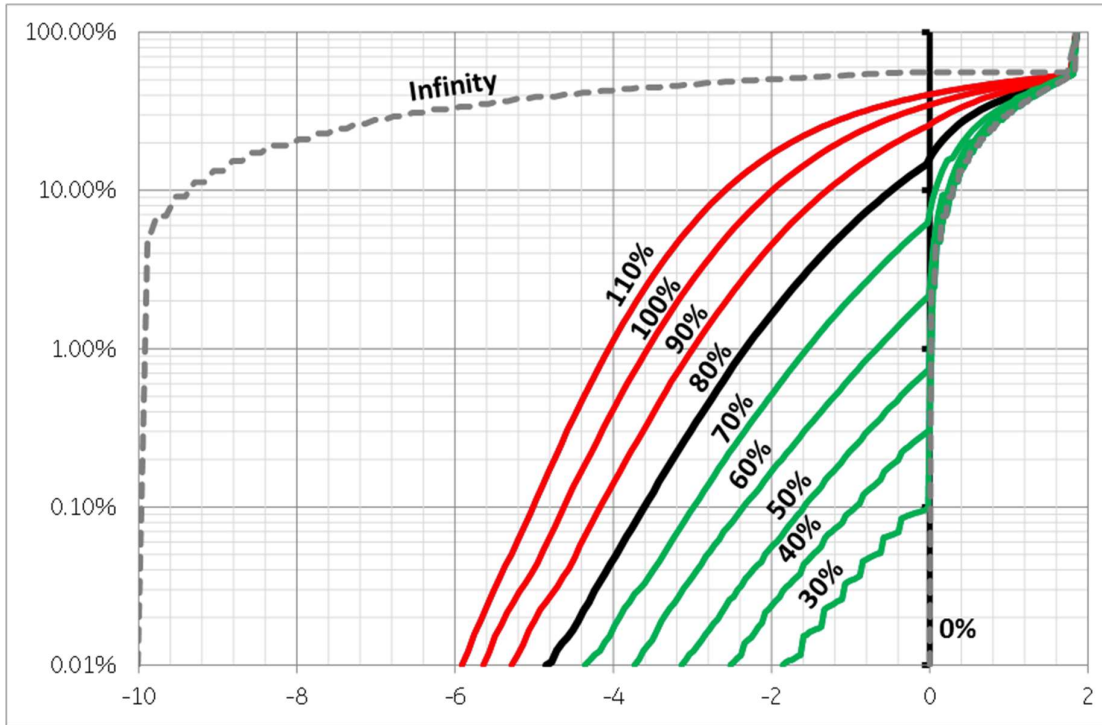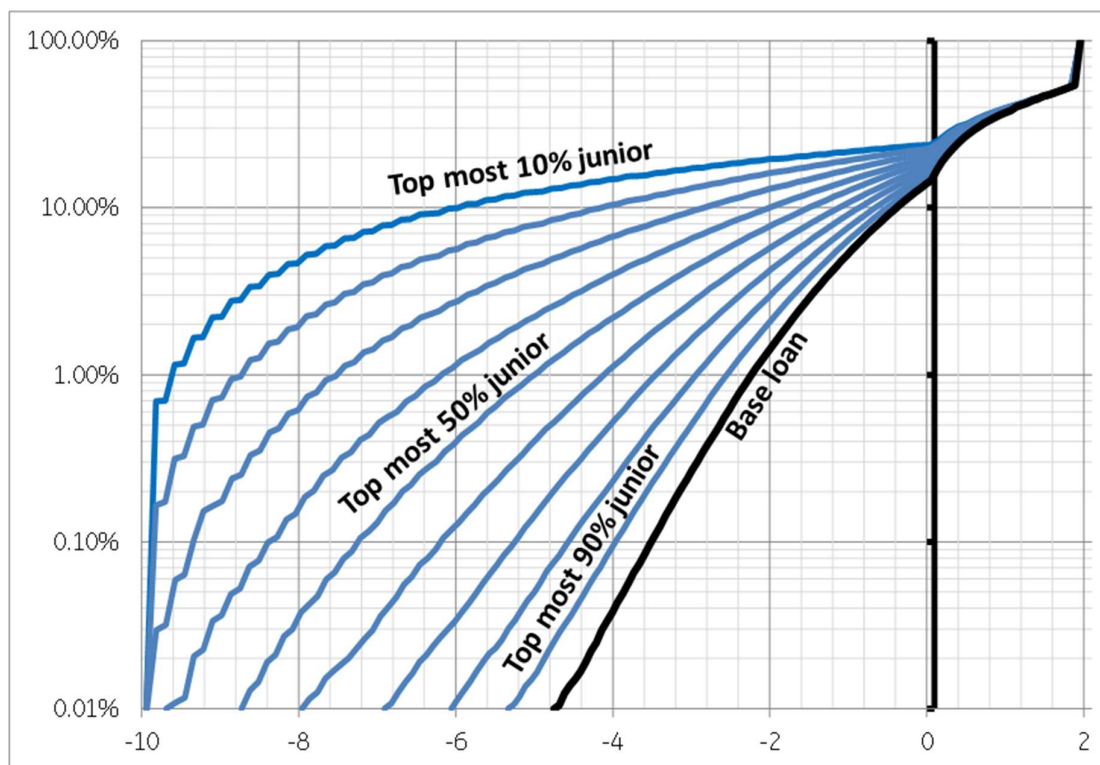*Figure 79: The Unit Loan with varying degrees of sub-ordination. The base loan is a senior loan. The "Top most 90% junior" means that loan represents 90% of the total exposure (senior+junior) and is subordinated to a senior loan that represents 10% of the total exposure.*

So, for example, the "Top most 10% junior" is a junior loan with an 80% LTV, but representing only 1.25mm of the 10mm loan where 8.75mm is senior ranking. The results have then been scaled up to 10mm to allow comparison between the different levels of subordination. Please note that the LTV in all cases starts at 80%. It is clear that the concavity of the NPV distribution, or the fatness of the tail, is more pronounced the higher the subordination level is.

This demonstrates why a VaR metric is not all that relevant for a single loan. Yes, the confidence level that we will not lose more than $10mm may well be 99.99%, (a AAA rating), but it is also 99.07%!

The Case for Asset Backed Finance

Loans and Structured Finance Products can be issued a "transaction rating" by a rating agency. While the corporate rating reflects the probability of default, the transaction rating reflects either probability of loss or probability and magnitude of loss, depending on which rating agency you ask. The idea of Asset Backed Finance is that the collateral will reduce the probability and magnitude of loss compared to an unsecured financing. In this way, a weak credit could borrow at lower spreads by offering collateral cover to the lender. In the following, we have translated the corporate rating to a probability and magnitude of loss for an unsecured exposure and compared to one with an aircraft as collateral, and then calculated the equivalent transaction rating. The Transaction rating is calculated based on the transaction capital as explained elsewhere in these notes using Monte Carlo simulation.
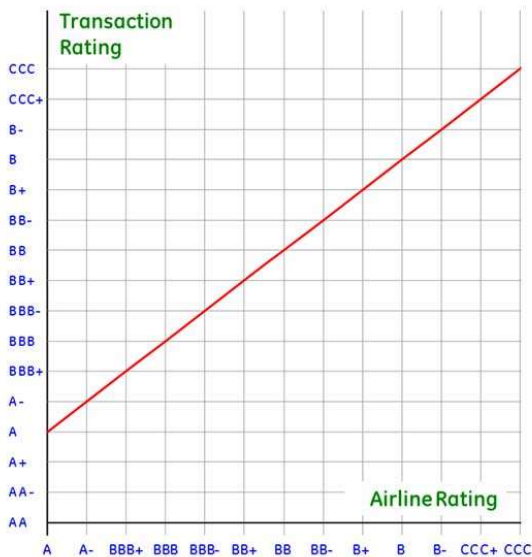
The picture to the left shows the Airline Rating on the horizontal axis and the unsecured Transaction Rating on the vertical axis for an **unsecured loan**. The loan is amortizing with a balloon payment that is recourse to the obligor. Here, the two ratings are identical as shown by the red line.
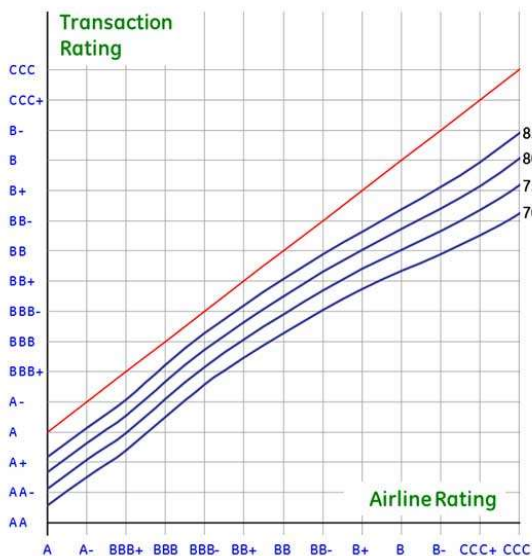
If we now look at a loan collateralized by an aircraft, in this case a B777-200ER, and calculate a Transaction Rating with different initial Loan to Value ratios, we see that the Transaction Rating is better than the corresponding Airline Rating by several notches. The effect is more pronounced the weaker the Airline Rating is.
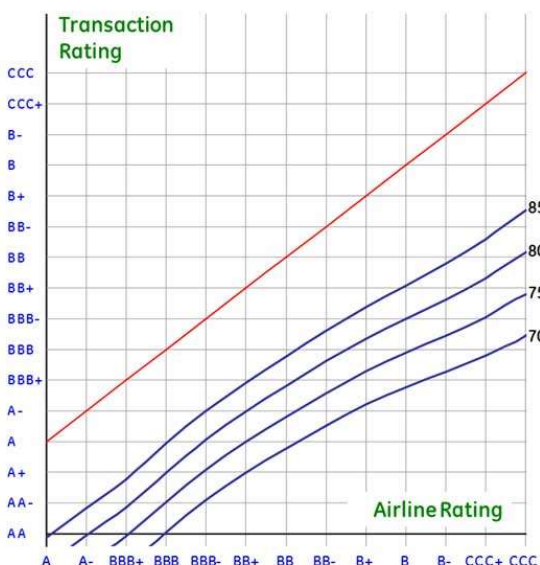
Switching collateral to a B737-800, we see an even stronger positive effect on the Transaction rating. The reason is that we have assumed that the B737-800 value is less volatile and that the switching costs for a B737-800 are lower than for a wide body B777-200ER. The B737-800 also has very good market penetration with many aircraft delivered to many different operators without too much concentration.

If we now make the balloon payment non-recourse (we cannot look to the obligor for repayment, only to the value of the aircraft), the Transaction rating is still better than the Airline Rating for weaker credits, but the reverse is true if the Airline rating is strong. In this case, the loss risk from the non-recourse balloon position dominates.

93

Portfolios

Up to now, we have looked at loans with a single collateral aircraft and a single obligor. When we build a portfolio of many different leases or loans, to different airlines and backed by different aircraft, we cannot simply add up the NPV distributions from each asset and expect to get the distribution for the portfolio. Defaults and aircraft values are not entirely dependent of one another. Many loans have multiple aircraft as collateral. Also, loans could be made to a special purpose vehicle that has many different loans to different airlines as security (Asset Backed Securities), or leases to different airlines.

It is unlikely that all aircraft values will move in exactly the same direction or the with same magnitude, or that all obligors will default at the same time. It is also unlikely that the value of an aircraft is fully dependent of the value of another aircraft, or that the default of one airline is totally dependent of the risk of default of another. In reality, there is less than full co-dependency or correlation. Correlation between two random variables (the value of two aircraft, or the default propensity of two airlines) is illustrated in the chart below. From left to right, the correlation is 0%, 40%, 80% and 100%.
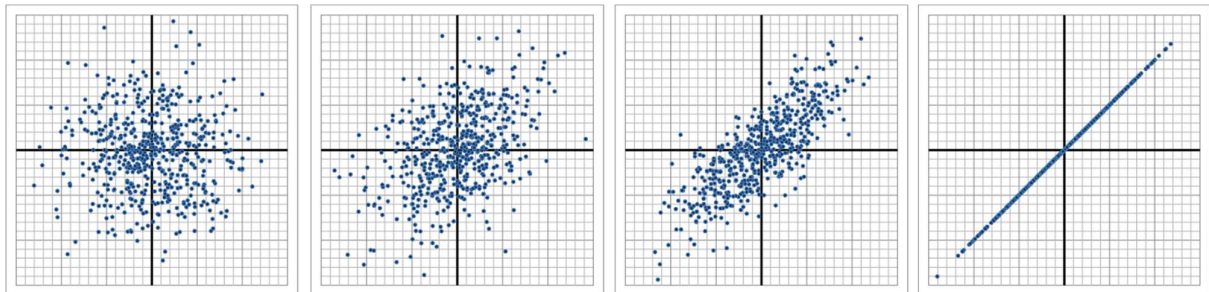
Figure 80: Illustration of correlation between two random variables, 0%, 40%, 80% and 100%.

It would be difficult in practice to measure the correlation between the values of two aircraft or the default risk of two airlines. And even if we could, the correlation of the past is not necessarily a predictor of the future. When we look at a portfolio with say, 100 aircraft, we would need the pair-wise correlation for 100 x 100 aircraft = 10,000 correlation numbers. This is an impossible task. One approach that is sometimes used in portfolio risk management is to use copulas, a statistical tool that defines the correlations between many variables. In the following, we will instead look at modeling the correlation and co-dependencies by using common influences. The objective is to figure out how aircraft values, airline default risk, and interest rates are influenced by common factors. The next figure (81) is an attempt to look at events and trends, and economic responses that would influence aircraft values. The principal value in an influence map like this is the benefit of going through the exercise of putting it together! It is a schematic, an opportunity to bring expertise around a flip chart and think about the world we live in and what could affect the value of aircraft.
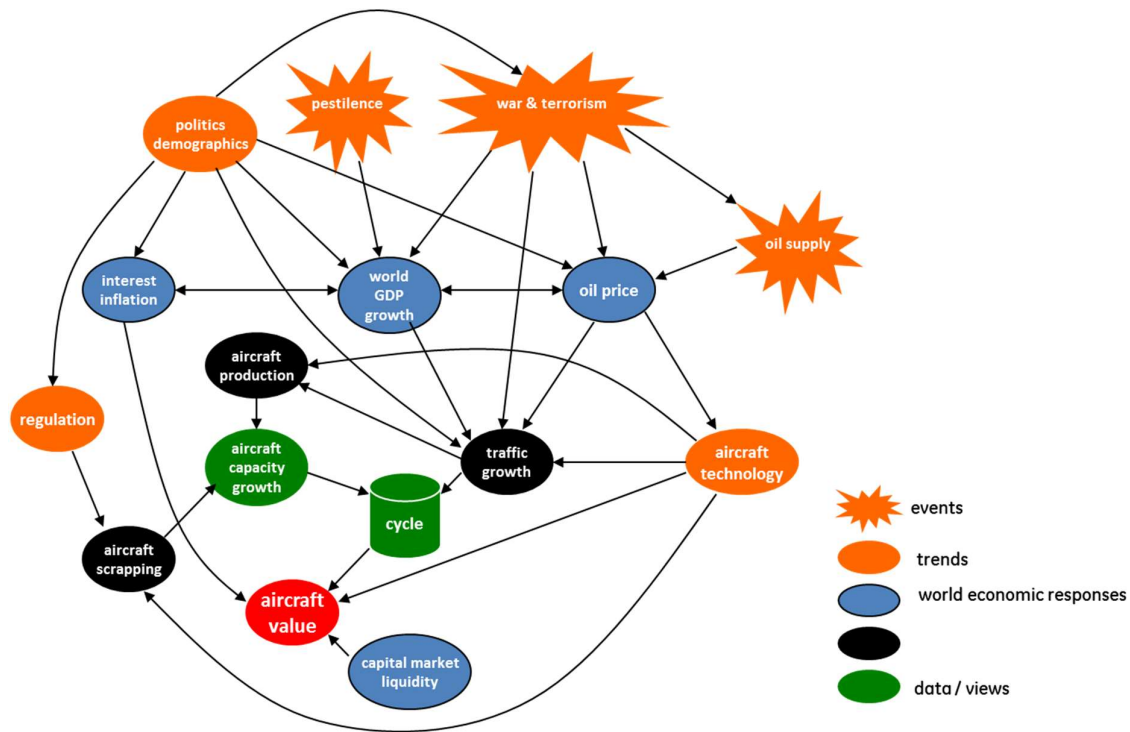
Figure 81: A simple influence map illustrating events and trends that could affect aircraft values.

We can then try to use this to create a more formal algorithm or model of the random processes for aircraft values, obligor default and interest rates. Figure 82 shows an example where the aircraft value is influenced by random variables that are global (same for all aircraft), group related (same for a family of aircraft), type related (same for a type of aircraft), or individual (unique for each individual aircraft). Maintenance value, as an example, is almost entirely derived from the individual source, since the maintenance status should be unrelated to the status of any other aircraft. Inflation, on the other hand, is global, and affects all aircraft the same way. Obligor default risk takes influences from a global source (world economy), a regional source (regional economy), and the individual or idiosyncratic source (unique to that one airline). In addition, the PURCS cycle affects everything as we have seen in the past, and also drives correlation between aircraft values, default risk and interest rates. This is a relatively simple extension of the single aircraft /obligor modeling we have discussed in the early chapters. It will require the use of Monte Carlo simulation. The thorny question is what weights should be applied to the various sources of random variables to replicate the co-dependencies that occur can be expected to occur in reality. This is a field of research that needs further exploration.
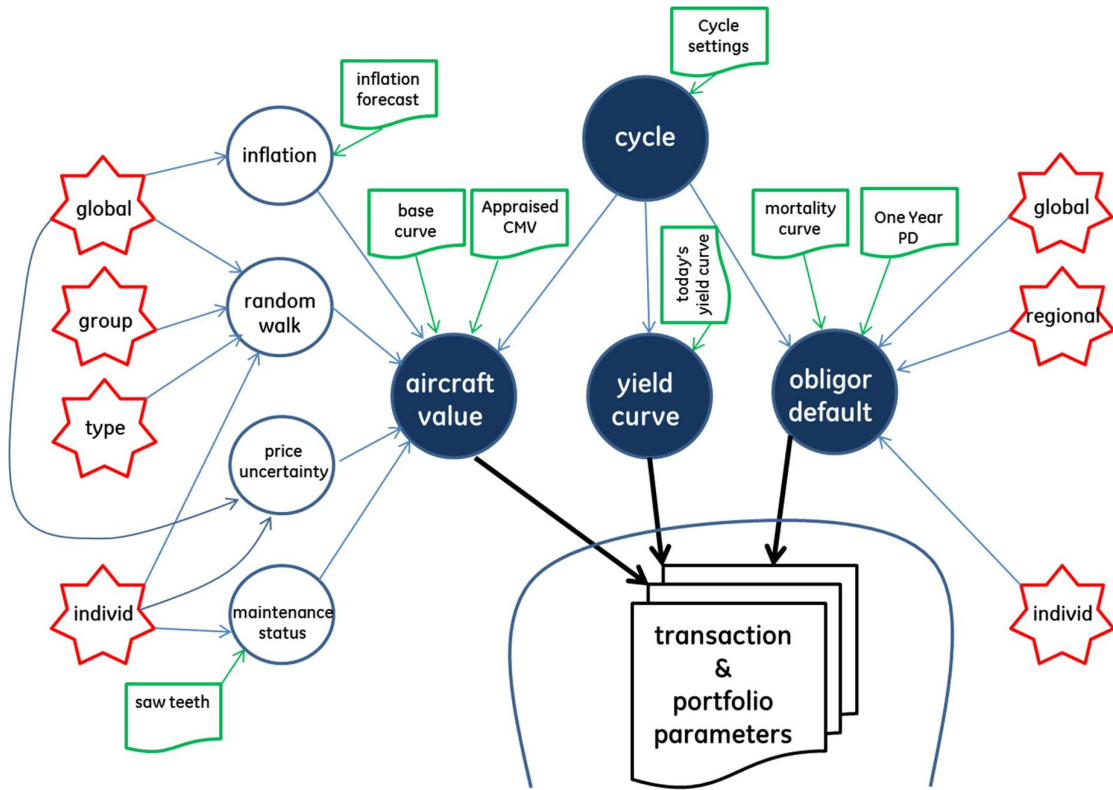
Figure 82: A schematic of joint influences to model co-dependencies of parameters in a portfolio of aircraft loans.

Diversification

The primary cause for portfolio risk being less than the sum of the risks of the consituent loans is diversification. The simple fact that all airlines are unlikely to default at the same time, or that aircraft values move in exactly in the same way, provides risk relief when assets are combined in a portfolio.
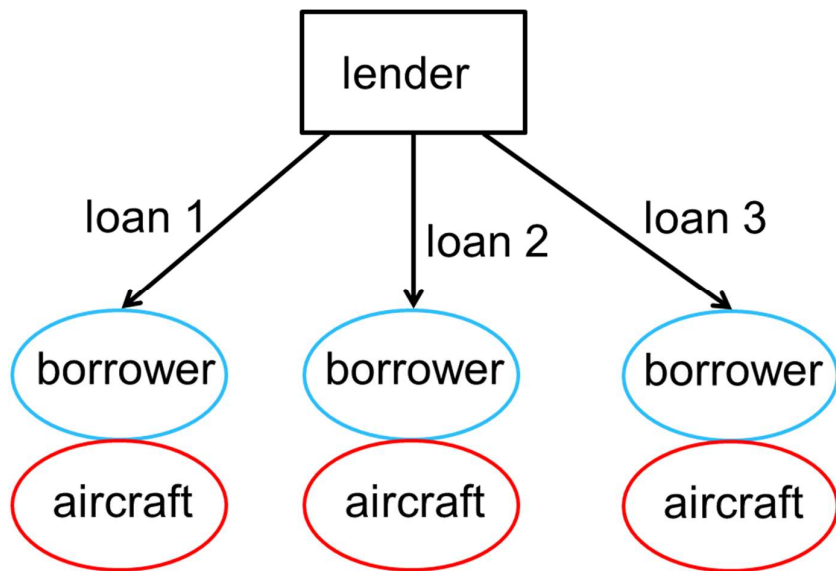
Figure 83: A portfolio of three unconnected Unit Loans.

The chart below shows the effect of diversification. The portfolio here consists of 64 1/64ths of Unit Loans with different obligors in different regions and aircraft of different groups and types. Every obligor has the same rating, every aircraft has the same expected value path, and every loan has the same term structure.
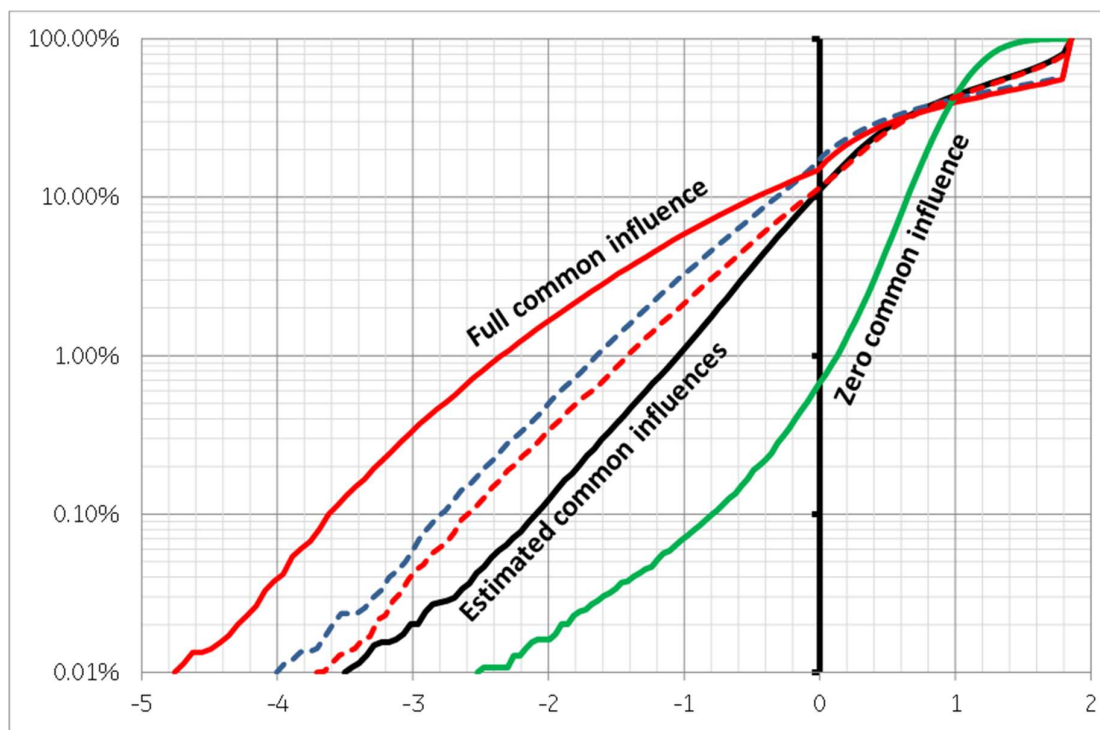
Figure 84: NPV-Probability distribution for a diversified portfolio of aircraft loans.

The effect of diversification depends on how strongly the default events and aircraft values correlate. The correlation in this example is caused by using common random variables to generate the aircraft values paths and default events. For aircraft, the base value path is influenced by a random variable that is based on contributions from four random variables: 1) the "World", 2) the "Aircraft Family", 3) the "Aircraft Type", and 4) the individual aircraft. The PURCS cycle is common in all scenarios while the maintenance status and instantaneous price uncertainty are idiosyncratic (zero common influence). For obligor defaults, the random variable takes contributions from three sources: 1) "World", 2) "Region" and 3) Individual. When the weight is set to 100% for "World" and zero for the remaining "buckets", we get the red line, this means that there is a very high degree of correlation in the portfolio. When "Individual" is set to 100% and the remaining weights set to zero, we get the green line. This is the case when we have almost independent obligor defaults and aircraft values. The correct weights are elusive, but the black line shows a setting where "World" is 45%, "Family" 20%, "Type" 25%, and Individual (Aircraft) 10%, while "Region" is set to 20% and "Individual" (Obligor) is set to 35%. The dotted blue line shows the NPV distribution when the aircraft are diversified but the obligors are not (single obligor – many aircraft). The red dotted line is the case where the obligors are diversified but the aircraft are not (many obligors – one aircraft)! Figure that one out! This is the kind of though experiment one can engage in in the ivory tower of a Monte Carlo simulator! Note though that the very best outcome is worse with many obligors, the chances are that some obligors will default before the end of the tenor, and the loans will earn less money. The effect of diversification is favorable on the loss side of the distribution. A gambler could take one big bet, and have a chance of a big pay-off, but also stand to lose a lot. A more prudent person would diversify, make several different and smaller bets, but of course, the big pay-off is less likely. On the other hand, the downside risks are reduced. Note that the loss reduction is significant at all percentiles of the loss side of the

distribution. The less correlated the assets (less common influences), the more the risks are mitigated.

The benefits of diversification are important but it is not something that will increase the market value of the individual loans of leases. The benefits arise from the way the underlying obligor and aircraft are less than perfectly co-dependent with other assets in the portfolio. A buyer of a loan or lease will not pay a premium for diversification since the link to the seller's portfolio is broken and the buyer, depending on its prior portfolio composition, will face different diversification dynamics. On the other hand, if the different aircraft or obligors are contractually linked, as in a cross-collateralized loan, there will be an improvement in the value of the loan.

Cross-Collateralization

Cross-collateralization means that the collateral from one loan also secures other loans. In its simplest form, we would have one loan secured by several aircraft in a pool. When the borrower defaults, the lender has recourse to all the aircraft to secure repayment. Obviously, the borrower would want some flexibility to dispose of aircraft and partially pre-pay, so in practice, the loan is often split in tranches, one for each aircraft. To safeguard the benefit of cross-collateralization, the lender would require some kind of "overage" or supplemental payment in excess of the outstanding amount that is being pre-paid for the tranche in question. The same would apply in case of a Total Loss and prepayment from insurance proceeds. The figure below shows a structure for a single obligor – multiple aircraft loan.
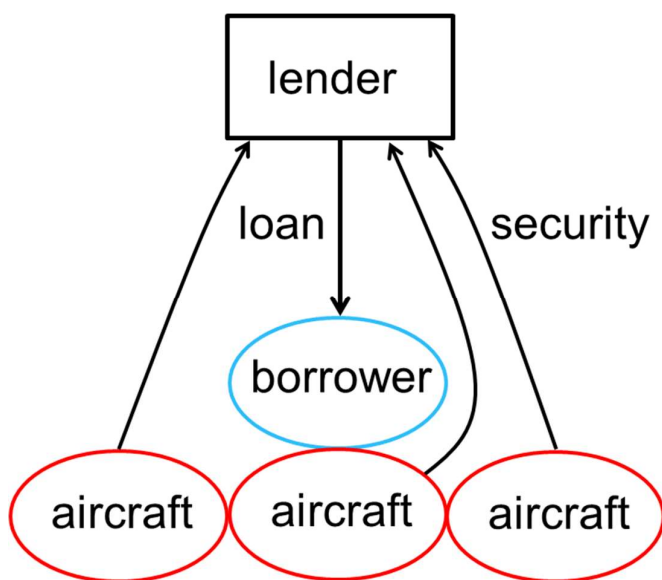
Figure 85: A structure chart for a single obligor – multiple aircraft loan where the aircraft are cross-collateralized.

A variation on this structure is that the lender has granted the borrower several different loans, each with one aircraft as collateral. The loan documents would then provide for the terms and conditions of the cross-collateralization rights. Sometimes, the "overage" payments could be contingent on loan-to-value covenants.

A next step is to have cross-collateralization within a pool of assets with different aircraft *and* different airlines. The flow chart below shows a structure for such a loan. To avoid the complexities of serially coupled obligors, this example shows a special purpose company as borrower. The key to cross benefits is that the covering collateral has "upside". An aircraft on

lease (operating lease) has potential upside. A loan has no "upside", it just has to be repaid, so to extract cross benefits from a pool of loans they need to be senior ranking with all the junior exposures *cross-subordinated*.
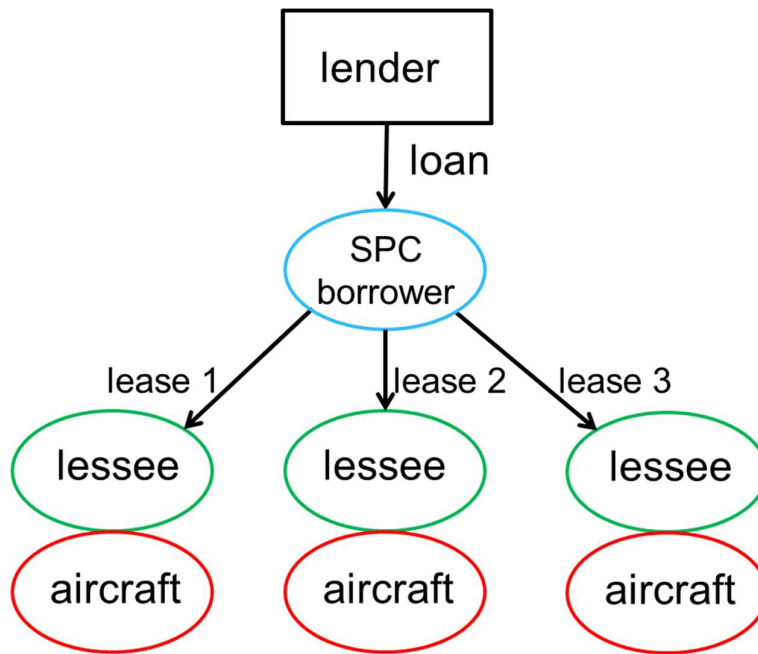
Figure 86: An Asset Backed Securities type of loan, with cross benefits from multiple obligors and multiple aircraft.

The chart below illustrates the effect of the cross-collateralization on the NPV-Probability distribution.

Figure 87: NPV-Probability distribution illustrating the effects of diversification and cross-collateralization.

The Black line represents the Unit Loan, shown here as a reference. The blue line shows the effect of diversification as the outcome of 4 Unit Loans with 4 different obligors, each with one aircraft but of different type and family from one another, run as a portfolio.

The green line shows the same portfolio but the four loans fully cross-collateralized.

The conclusion we can draw from this chart is that diversification significantly reduces the loss risk at all levels but also limits the probability that we will achieve the maximum possible Net Present Value. This is very intuitive – the chance that all of the four loans will be default free and loss free in the same scenario is more remote than that the single Unit Loan will.

Cross-collateralization on the other hand fully dominates the non-crossed portfolio, but for extreme losses or very good outcomes, the cross makes no difference. Thus, cross-collateralization is more of an Expected Loss mitigant than a VaR mitigant.

The beneficial effect of cross-collateralization is rather modest in the example above. We need to bear in mind that in this portfolio of Unit Loans, all starting LTVs are identical, and all obligor ratings are the same (which does not mean that they will all default at the same time in each scenario, only that they are all expected to default with the same frequency *on average* in multiple scenarios).

Figure 88: NPV-Probability distribution for a portfolio of two loans with varying inequality between LTVs.
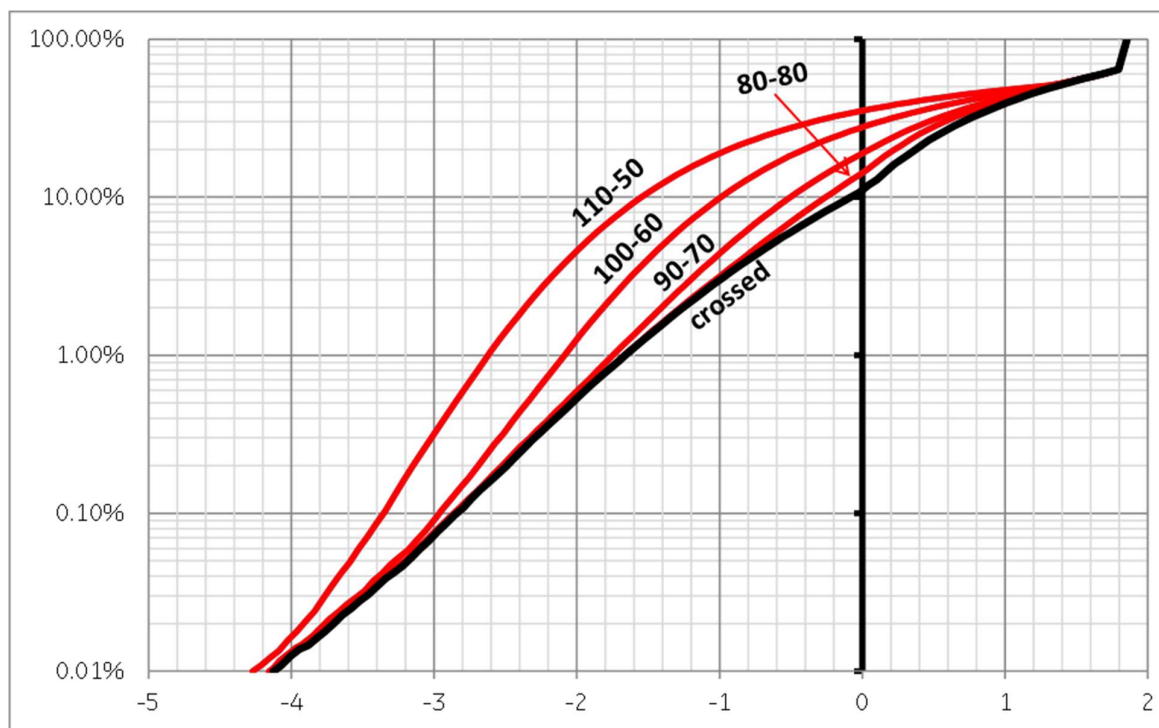
When the loans in the cross-pool are uneven with respect to LTVs, the benefit of the cross provision increases. The chart above shows the effect of cross-collateralization between two unit loans when the LTVs differ. The LTVs have been shifted pro rata over the term structure. The red lines show the NPV distribution for two non-crossed loans at LTVs 110%-50%, 100%-60%, 90%-70% and 80%-80% respectively. The black line shows two unit deals with full cross. The outcome is identical to this case for any LTV combination as long as we have full cross provisions. The take-away here is that the full cross eliminates the adverse consequence of having uneven LTVs. In other words, the more uneven the LTVs are, the more valuable the cross-collateralization is for the lender. The benefit vanishes for VaR at very high confidence levels (when things get really, really bad, there is no escape – even low LTV loans will be "under water").

Note that for Operating Leases, the whole concept of cross-collateralization is moot. The lessor already "owns" all the upside in the assets, and therefore a cross-collateralization would not make any difference. Diversification, on the other hand, is very beneficial, just like in a loan portfolio.

Cross-Collateralization Benefit on a Single Loan

We have discussed how diversification is beneficial, but does not add market-value to a loan or lease, and how cross-collateralization does. Let us now take an example of a lender that granted a loan to a borrower for a new aircraft, say, five years ago. The loan has been partially paid down and the LTV is now lower than it was at the inception of the loan. The borrower is taking delivery of another new aircraft and is seeking a new loan from the same lender. In order to convince the lender to advance a high loan amount, the borrower offers the lender cross collateralization from the existing loan. If we analyze the two loans with cross, we will inevitably count the diversification benefit arising from the fact that we have two different aircraft. Even though the obligor is one and the same, and the aircraft type in this example is

101

the same, there is still some diversification benefits linked to the maintenance value and instantaneous price uncertainty that are idiosyncratic risks rather than systemic. So how do we account for the cross benefit in the new loan without accounting for the diversification?

The charts below show the existing loan (to the left) and the proposed new loan (to the right).
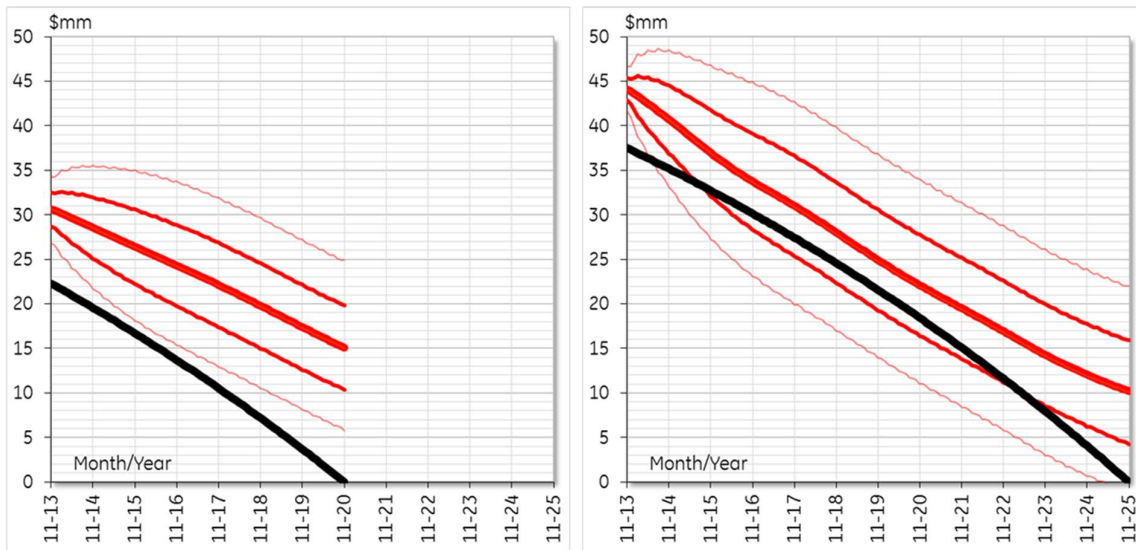
Figure 89: The exposure is shown in the black line, and the projected aircraft value is shown in the red thick line, with one and two standard deviations up and down.

The obligor is assumed to have a 10% one-year default probability. The chart below shows the NPV-Probability distribution for the existing loan (the gray line) with a current exposure of $22.3mm, a LTV of 73% and a Risk-Reward Ratio of 2.2%. The proposed new loan (black line) has an initial loan amount of $37.5mm, a LTV of 85% and a very poor RRR of 28% - something the lender would not accept.
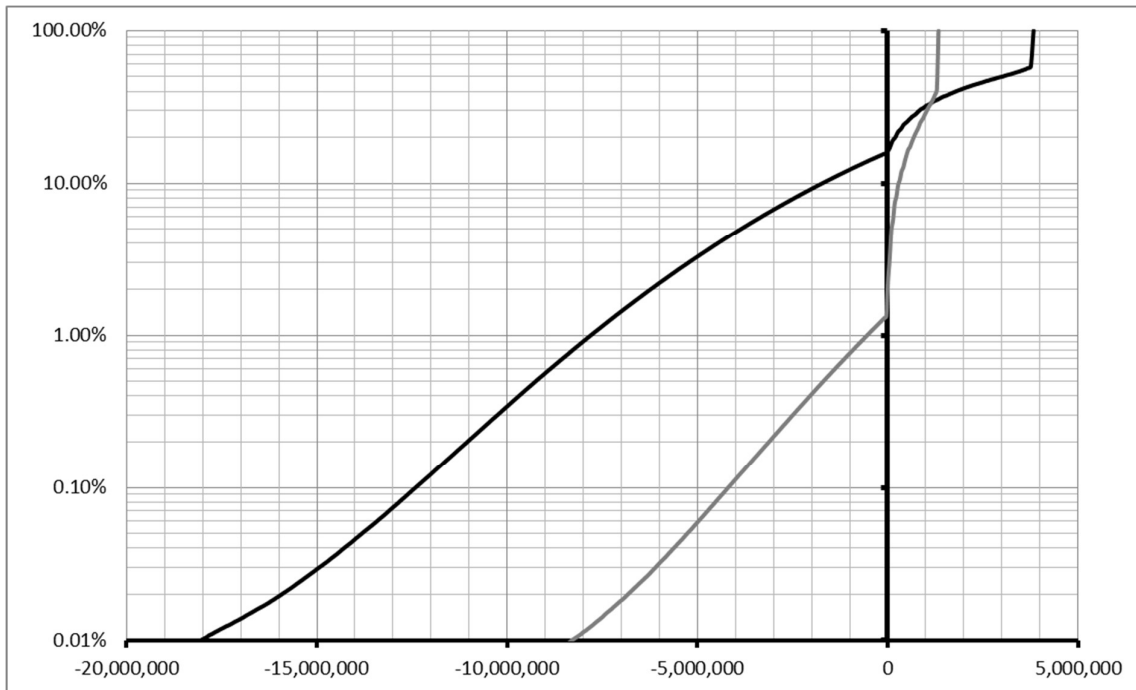
Figure 90: NPV-Probability distribution for an existing good loan and a proposed new loan, unattractive on a stand-alone basis.

102

The cross provisions are agreed to be full (all sales proceeds from the aircraft securing the existing loan, once the existing loan has been fully repaid, goes to the waterfall of the new loan. However, when the existing loan reaches maturity, the mortgage over the aircraft will be released. In this sense, the cross is limited in tenor). To estimate the NPV-Probability distribution for the new loan only, but with the cross benefits, we proceed as follows. We build a portfolio of the existing loan and the new loan with the cross provisions modeled, but then *deduct* the existing loan *on a stand-alone basis*. In fact, Monte Carlo simulation gives great opportunities to build complex transactions with smaller building blocks.

The NPV-Probability distribution for the new loan, taking the cross, but not the diversification, into account (green line) is shown in the chart below. The RRR is 12%, significantly improved by the cross. Again, note that the maximum earnings are obviously unchanged, and that the loan with cross dominates the non-crossed loan all the way until the loss tail reaches around a 99.9% confidence level.
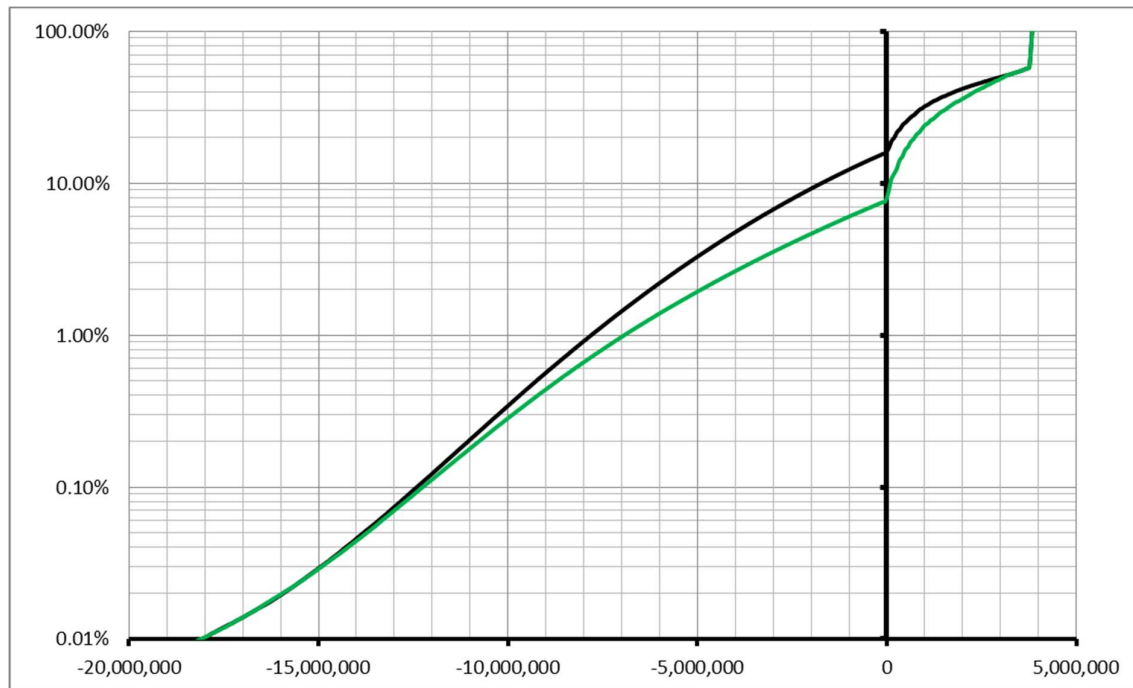
Figure 91:NPV-Probability distribution of the new loan non-crossed (black line) and crossed (green line).

From a mark-to-model perspective, if we assume a market RRR of 20%, the cross benefit is the equivalent of either having a $2.1mm cash deposit for the full term as additional security, or adding a cash position of $0.94mm to the transaction. The chart below illustrates this. The black line is the new loan without cross benefit, the green line is with cross benefits, the red line shows the new loan without cross but with a $2.1mm cash deposit (note that the red line is shifted $2.1mm to the right of the black line on the loss side of the distribution), and finally, the blue line shows the new loan without cross but with a $0.94mm cash position added (note that the blue line is shifted $0.94mm to the right of the black line for the *entire* NPV spectrum). The green, red and blue lines represent the same mark-to-model value, but the risk-reward properties are not identical. The cash deposit has the least extreme risk, the crossed loan has the least risk when the confidence level is below 95%, and the loan with the cash position carries more risk than the crossed loan and the loan with a cash deposit below the 99% confidence level, but has higher reward potential (the lender gets to keep the cash in all cases!).

Figure 92: The new loan non-crossed (black line), crossed (green line), with a $2.1mm cash deposit (red line), and with a $0.94mm cash position.

Without this type of probabilistic analysis, it is very difficult to estimate the benefit is cross-collateralization.

A naïve approach to estimating the benefit of the cross would be to say that we have a certain amount of cushion between the contractual exposure and the projected aircraft value under the existing loan, a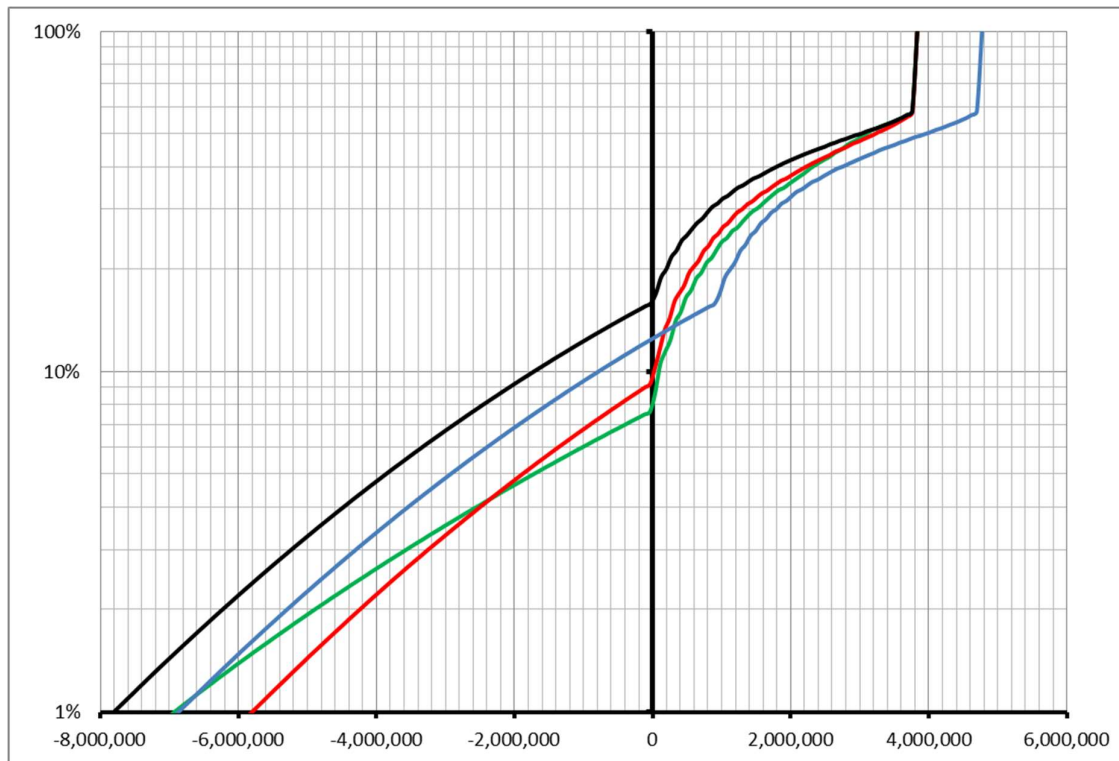nd to look at this as additional collateral on a purely static basis, like a cash deposit. The RRR in this case would be 1.8%. While the reward side of the distribution is basically unchanged, the mitigation of losses is greatly exaggerated. The fallacy, of course, is that the collateral cushion from the existing loan is very fragile. Not only is the protection derived from a value that is uncertain, but it emanates from the "junior" portion of the value, available to cover the new loan only when the "senior" exposure of the existing loan has been fully extinguished.

Portfolio Capital and Transaction Capital

The traditional approach to pricing a loan or a lease is to get compensation for the expected loss rate, transaction and handling costs and then achieve a target return on the risk capital. Too often, the risk capital is based on a set leverage ratio, the same across a whole product type. The risk of mispricing is obvious as some transactions carry more risk than others, and the expected loss rate can only compensate for a part of the higher risk. It would be better to use a transaction leverage that somehow reflects the riskiness of the loan or lease. However, we have seen that the aggregate VaR from individual transactions is far higher than the VaR for the full portfolio, due to diversification and cross-collateralization benefits. So how could the portfolio VaR be disaggregated into individual transactions? The shape of the loss tail is different for different transactions. The VaR is a poor metric to capture the bad case downside risk in a single transaction, and the two following charts illustrate why.

The chart below shows the NPV-Probability distribution for a loan where the VaR for different confidence or rating levels is well separated. For the BB rating, the VaR is a positive NPV, indicating that the capital would be negative! The BB level is simply too low a confidence level for purposes of setting risk capital.
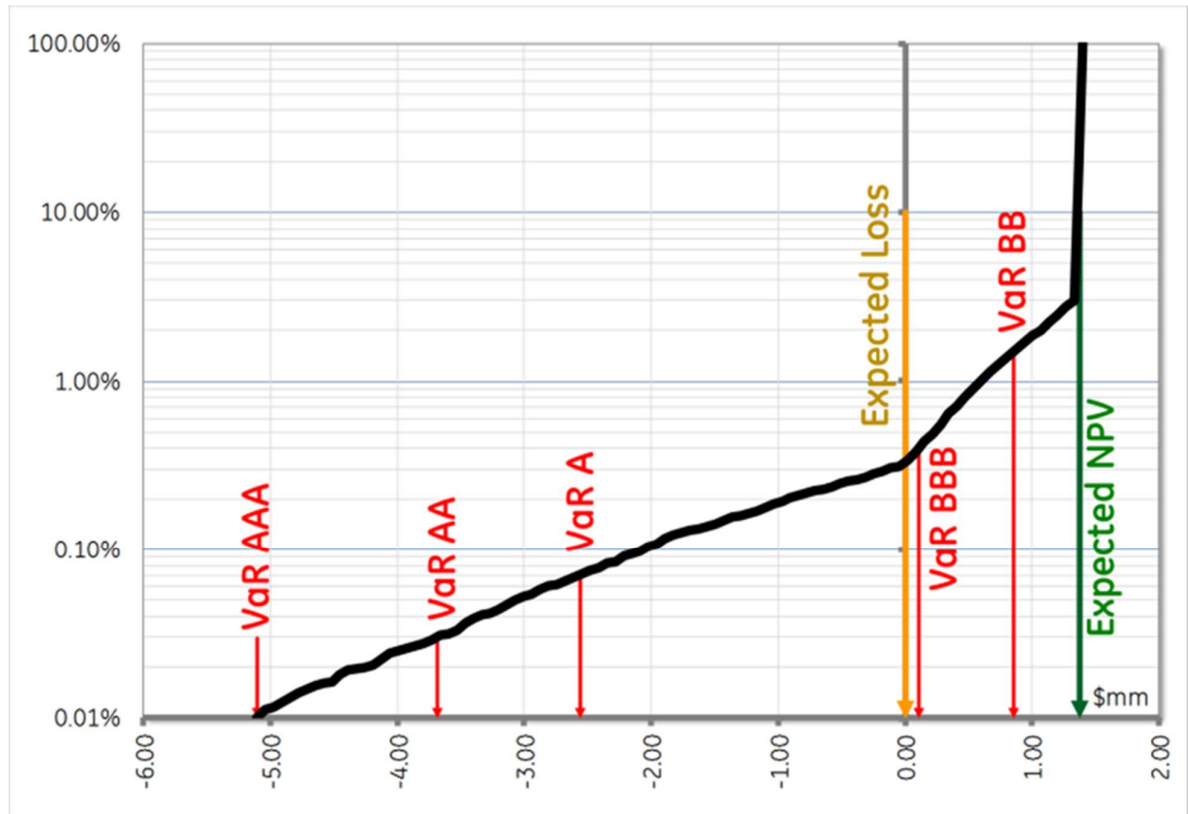

Figure 93. NPV-Probability distribution for a loan with VaR levels marked out for different confidence or rating levels.

Now look at another loan, a really bad one, with a really fat loss tail. The VaR for the BB level and AAA levels are almost the same, there is very little differentiation in VaR. This further illustrates the weakness of VaR on a transaction level.
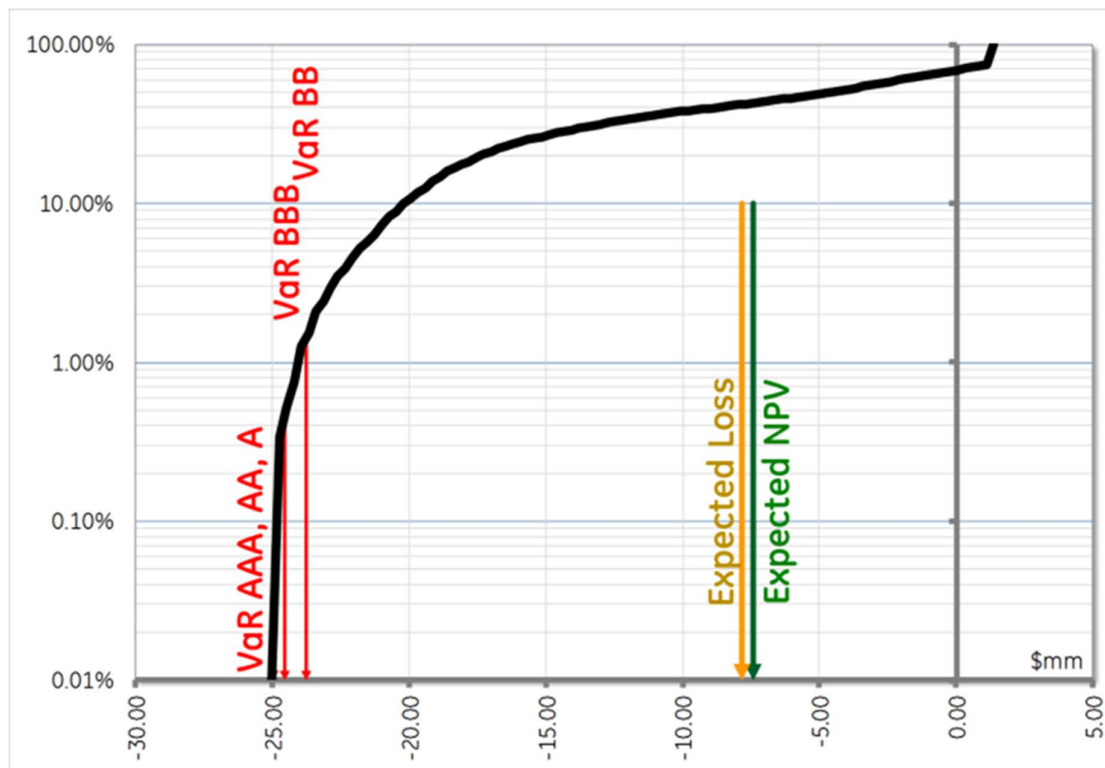
Figure 94. NPV-Probability distribution for a loan with a fat tail. The VaR numbers are closely clustered around the same number, a loss of all the principal.

The idea with Capital is that one would set an amount aside to be certain with a given degree of confidence that the loss will not exceed the capital and wipe out the transaction/portfolio. This concept is really meant for a portfolio rather than for a single loan. Nevertheless, it is practical to assign some kind of "unexpected loss" to a single loan. The approach of calculating the Value at Risk from the NPV distribution has several limitations as seen above. One being that for a loan with a very fat tail (such as a junior loan, where the entire exposure may be wiped out in an adverse event), the loss at a given confidence level may be the same as that at a much lower confidence level. Another limitation is that for a very thin loss tails, the VaR may be a positive number at low confidence levels (negative capital). Also, the portfolio VaR depends, among other things, on concentrations in the portfolio. This would be lost if one were to aggregate the transaction capital for a whole portfolio. There is probably no way to disaggregate the portfolio capital and allocate to the constituent parts of the portfolio correctly, but a pragmatic way could be as follows: The VaR is calculated for five different confidence levels for the entire portfolio (by simulating the NPV for the entire portfolio in each scenario). The stand-alone VaR is then calculated for each individual loan at the same five confidence levels. The aggregate VaR for each confidence level is calculated and compared to the corresponding portfolio level VaR. This comparison is used to calculate a scaling factor for each confidence level. These factors are then used to re-scale the VaR amounts for each individual loan and for each confidence level. The confidence level that yields the maximum (largest negative NPV) re-scaled VaR is then used to pick the re-scaled VaR for each transaction. Finally, the aggregate of this re-scaled VaRs is compared to the single A level portfolio VaR and this yields a diversification factor which further reduces the adjusted transaction VaR. From this amount, the Expected Loss of the stand-alone deal is deducted. We finally set a floor of 0.5% of the exposure for a given transaction and call this Transaction Capital. With this process, the individual NPV distributions are compared to the shape of the full portfolio distribution, and the total capital depends not only of the constituent loans/leases, but also on the amount of name- and aircraft type-concentration in the portfolio. The drawback is that while the concentration issues are reflected in the overall level of capital, the contribution of each individual transaction to this concentration is lost.

It could be meaningful to calculate the *marginal transaction capital*, what is the incremental VaR of my total portfolio when the transaction is added? This would depend not only on the properties of the transaction, but also of the content of the prior portfolio. The contribution to the concentration risk would be visible at the individual transaction level. Unfortunately, we have found that the resolution of the Monte Carlo simulation on the overall portfolio level is not fine enough to capture the marginal contribution from an individual deal with any meaningful precision.

On Fat Tails

It was mentioned on page 80 that the extreme end of the NPV-Probability distribution cannot be estimated. When we look at the Value at Risk, the loss that we think will not be exceeded with a given confidence, say 99%, we are thinking of events that would only happen once in a hundred trials. At the AAA rating level, we are looking at 99.99% confidence, an event that could be expected to occur once 10,000 trials. We obviously do not have any *relevant* data stretching that far back that could help us model loss events in that end of the probability spectrum. Going shorter would not give us *sufficient* data. So we are stuck. From the start of quantitative risk modeling, loss distributions were often assumed to be Normal or Gaussian. Stock movements in the Black & Scholes option pricing formula are assumed to follow this distribution. However, historical data of stock prices show that the empirical distribution has "fatter tails". Extreme movements occur more often than the Normal distribution would suggest. Nassim Taleb discussed this in his book "The Black Swan". Natural phenomena like earth quakes, rainfall, solar flares, or social phenomena like city size, wealth or book sales, more resemble a power law distribution than a Normal distribution. The power law in the tail of the NPV-Probability distribution would show up as a straight line in the logarithmic scale we have used for the probability axis in this section. The chart below shows an example of the NPV distribution in a $6bn aircraft backed loan portfolio with some 520 different aircraft and 90 different obligors.

In this example, 20 million scenarios were run (that took a couple of weeks!). The line is straight (on a log scale) below a probability of 0.01%. This is consistent with a ***power law distribution***. If we extrapolate the straight portion of the tail, it would intercept the negative $6bn mark at a probability of around 0.0001%, a one in 10 million event. Could one actually lose the full amount of the portfolio? Well, we could always imagine the invention of teleportation (which would instantly save us from check in lines, security controls and bland food in plastic boxes), or a massive asteroid striking earth, or the collapse of the Higgs field, obliterating the universe, although the empirical data from these events would not be all that useful for modeling our next aircraft loan or lease. What is clear though, is that any model that suggests a normal distribution in the tail end of the portfolio is most likely wishful thinking.
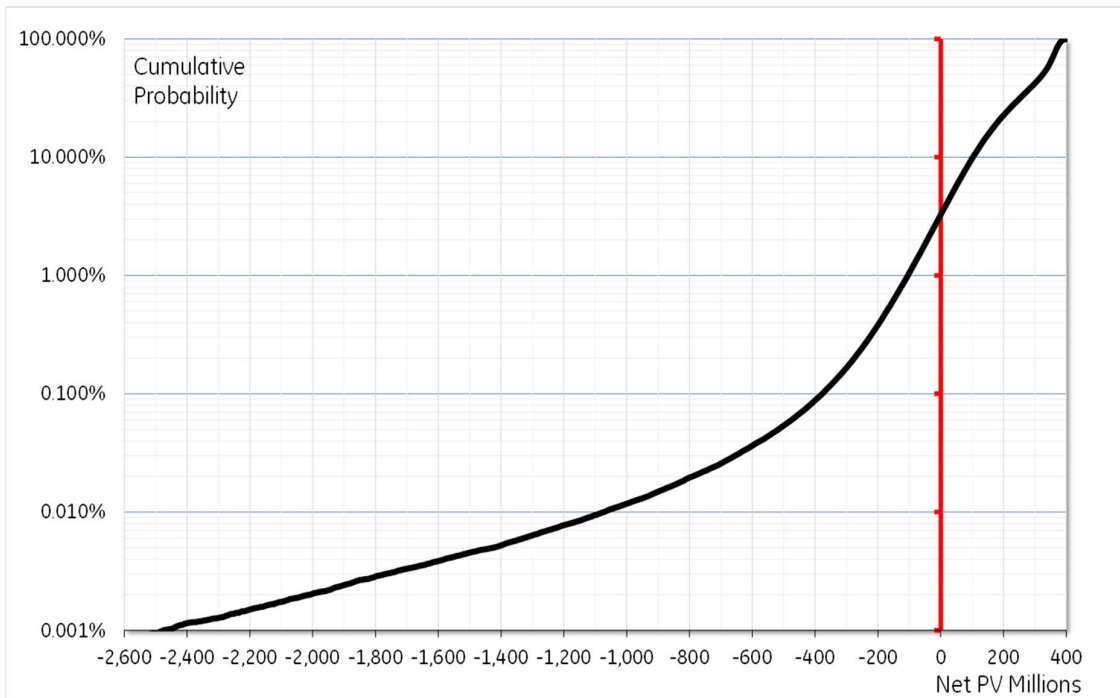
Figure 95: NPV-Probability distribution for a $6bn loan portfolio. In this example, the probability of loss is around 3%, maximum potential PV is around $380mm, VaR at the 99.9% confidence level is around -$380mm, at the 99.99% confidence level is around -$1.3bn, and at the 99.999% level around -$2.5bn. The shape of the tail has no support by empirical data, and will never have. The point though is that the tail is much fatter than a Normal distribution tail (more kurtosis).

Loan Prepayment Risk

In many loans, the borrower has the option to voluntarily prepay a loan before contractual maturity. There is no risk of "loss" for the lender in this case. However, the lender will forego earnings as well as getting rid of the risk. Since the option is for the borrower to exercise, prepayments typically happen when a loan is valuable for the lender and off-market for the borrower. Therefore, the prepayment option typically has a value to the borrower, and a corresponding negative value to the lender. This prepayment risk can be mitigated by prepayment fees or other restrictions on prepayment. The value of an early termination option can be modeled, although it is less straight forward than valuing the loan itself. There are no market values for prepayment options that can be observed so the calibration is difficult.

Prepayment Speed

Prepayment speed in a loan portfolio is the amount being prepaid during an accounting period divided by the starting balance of that period. The quantification of prepayment could also be based on the prepaid Balance Volume (amount x contractual duration representing the area under the principal repayment profile – this is a more meaningful measure for a lender). Another useful quantification of prepayment is the realization ratio – the ratio between the achieved average life (up to the prepayment date), and the contractual average life.
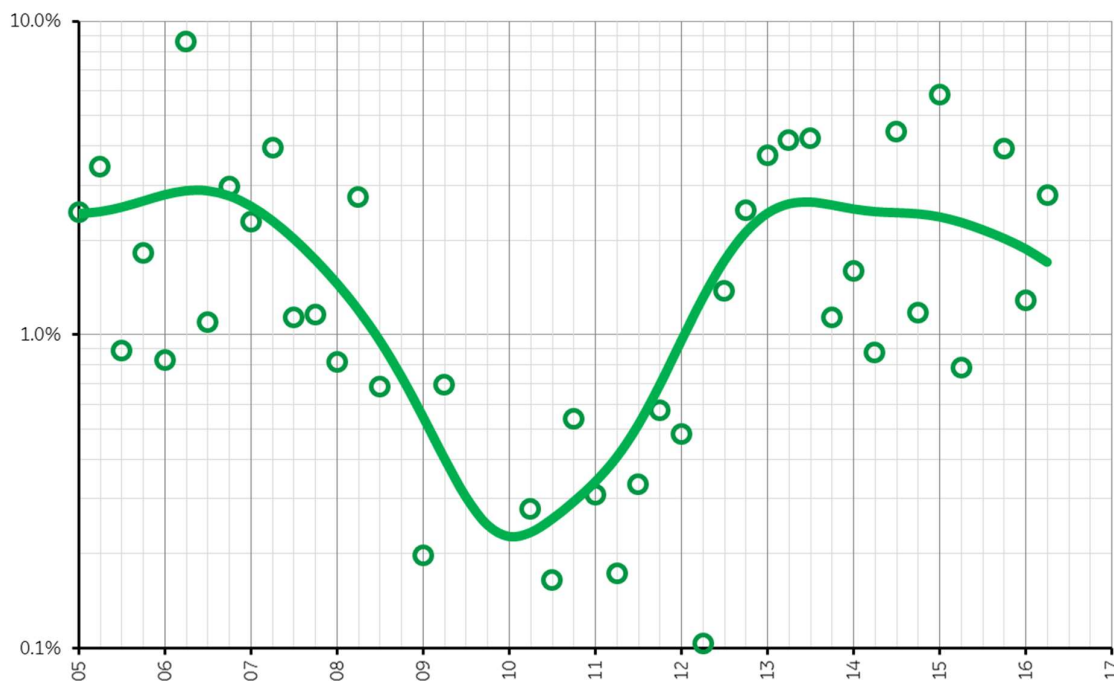
Figure 96: An example of quarterly prepayment speeds in an aircraft loan portfolio on a logarithmic scale. The dots represent the prepayment speed in the quarter and the line is a trendline through the scatter using a Wiener filter. Note how the prepayment speed slowed in the aftermath of the financial crisis 08-09.

Prepayment Incentive – The Three Cs

A borrower could decide to prepay a loan before contractual maturity for a number of reasons. Here is a list of some that come to mind:

Cost

1) The borrower may have a **cost incentive** to refinance. While one could assume that the terms and conditions were "in the market" when the loan was first negotiated, a number of factors may have "drifted" since the date of closing.

a) In fixed rate loans, the **base rate** may have decreased for the remaining duration of the loan. Unless the swap breakage cost is payable by the borrower upon prepayment, the borrower could take the opportunity to refinance at the lower market swap rate and thus save the present value of the difference between the interest payments calculated using the contractual rate and the market rate.

b) The borrower's market **credit spread** may have compressed thanks to i) an improved **credit rating** since the closing date, an improved **market perception of the collateral aircraft**, iii) an improved **loan to value ratio** from the principal run-off or an improvement in the aircraft value, iv) an improved **cyclical outlook** for airline credits or aircraft values, v) a higher **market Risk-Reward Ratio** due to an increased risk appetite among banks and investors. By prepaying and refinancing the loan, the borrower could now save the present value of the difference between the interest payments calculated using the contractual Margin and the Margin that could be obtained for the exposure in the current market based on the parameters above.

c) In loans with non-recourse balloons, a deterioration in the projected balloon LTV will obviously reduce the incentive for the borrower to prepay. A non-recourse balloon recovery shortfall is the lender's problem, and its corresponding cost is the borrower's upside.

Cash

2) There may be a **cash incentive** for the borrower to refinance. As the original loan amount amortizes faster than the aircraft value depreciates, the LTV ratio reduces compared to the closing date. The borrower may decide to use the collateral aircraft to extract more cash at secured lending terms.

Collateral

3) The borrower may want to **dispose of the aircraft**. If the aircraft becomes surplus to the airline, or if the borrower simply sees an opportunity to sell the aircraft on favorable terms, the loan will be prepaid and the mortgage lifted.

There are also disincentives to prepay. The borrower will incur transactions costs for refinancing a loan, and the prepayment option will typically entail some type of penalty of prepayment fee.

Incentives described above may turn to disincentives when the obligor rating drifts, the LTV, or any of the other parameters described above move the wrong way.

Why Model Prepayment Risk?

There are several reasons for trying to model prepayment risk.

1) To optimize Asset-Liability Management, a financial institution may seek to match the funding of the loan portfolio to expected duration rather than contractual duration.

2) To correctly value a loan, the mark-to-market value needs to include the value of the option to prepay. This valuation may also be required by accounting rules, and specifically when prepayment options are deemed to be "embedded derivatives".

3) Conditions of prepayment and prepayment fees need to be negotiated in the context of an overall optimization of the loan terms in a coherent and consistent manner. As we will see in the following, the optimal prepayment fee structure will depend on where we are in the industry cycle, and will also depend on the type of borrower.

4) Financial planning and budgeting improves if prepayment risk is accounted for.

When is the Prepayment Option Exercised?

A rational borrower should exercise the option to prepay when the remaining mark-to-market value of the loan (without prepayment option) plus the value of freeing up the collateral to be used for other financing purposes, exceeds the prepayment fee and transaction expenses to refinance. However, if the option is not exercised at a given point in time, the option remains, so it is a question of finding the optimal time to exercise. In this respect, the problem resembles the valuation of an American option (a call option that can be exercised at any point in time). Given the large number of risk factors in an aircraft backed loan, this becomes a very complex exercise. It is also very difficult to find a robust explanatory model for how the cash incentive affects the prepayment risk.

Empirical prepayment data shows that there is some degree of irrationality or uncertainty in prepayment behavior, or at least, some factors that influence prepayment cannot be reliably forecasted. The disposal of an aircraft is a case in point.

Regression Models

A simple approach to modeling prepayment risk is to perform regression analysis on historical data. We have to select the regression parameters and test for a series of historical dates, and look at both loans that were prepaid, and loans that were not prepaid. Thereafter, we determine which parameters or attributes of the loans best explain prepayment behavior. Here is a list of "suspects": 1) the yield curve, 2) high yield bond indexed or cycle amplitudes that would be a proxy for market credit spreads, 3) the loan LTV, 4) the borrower rating, 5) loan age, 6) prepayment fees, 7) loan type (airline loans, loans into tax advantaged structures would tend to prepay less frequently, whereas other investor loans typically prepay more frequently). The type of regression analysis that we apply here is called a multi-nominal logit regression where we have a combination of continuous variables (LTV, loan age, …) and binary variables (yes of no for investor loan, tax loan, ….). Independent variables are then selected using univariate regression (vary one parameter at a time) to see which ones best explain Prepayment Speed.

Monte Carlo Modeling

We have discussed Monte Carlo Modeling in the previous sections of these notes. If we have a Monte Carlo model for the loan, we can easily build in the prepayment aspect into this model. The beauty of Monte Carlo is that we generate distinct scenarios with respect to aircraft values, default and yield curve. We can then measure the propensity to prepay in each scenario.

Cost Incentive

1) We compare the simulated yield curve to the contractual base rate in the loan and can then calculate the cost saving (or increase) the borrower would obtain in a refinacing on a present value basis at each payment date from the analysis date until maturity. This calculation accounts for the way swap break cost and gains are contractually attributed in the loan. Very often, the swap break is passed through to the borrower, in which case the yield curve has very little impact on the cost.

2) We compare the "market loan margin" at each payment date to the contractual loan margin and calculate the present value of the cost saving at each payment date. The question is how we estimate the "market loan margin". If we assume that the contract margin represents the correct market margin for that specific borrower, for that specific aircraft and repayment profile at the ***closing date*** of the loan, we could assume that the drift in the amplitude of the PURCS cycle (which is simulated in our model) from closing could be used for modeling the "market margin". Loan spreads and PURCS correlate very well historically.

3) Since the obligor rating is not simulated per se in our model, we use the default point of the obligor in each scenario. We calculate the net present value of the cost saving (or increase) due to the drift in market margin from the analysis date until the *earlier of maturity and default*. Borrowers with a weak credit rating will experience a higher frequency of simulated scenarios where there is a default prior to maturity, and hence the calculated cost incentive will reduce for weak obligors.

4) For loans with a non-recourse balloon, in scenarios where the balloon amount exceeds the collateral value simulated at maturity, this shortfall will be present valued and deducted from the borrower's cost saving.

5) Any prepayment fees are accounted for in the present value calculation.

6) We assume the borrower will incur transaction costs to refinance the loan on the better "market" terms. We use the same transaction cost that we assumed in the model when the loan was first entered into.

We can now calculate a profile of the cost saving (or increase) over the remaining loan term that the borrower would face in the event of a prepayment, in each scenario of the simulation.
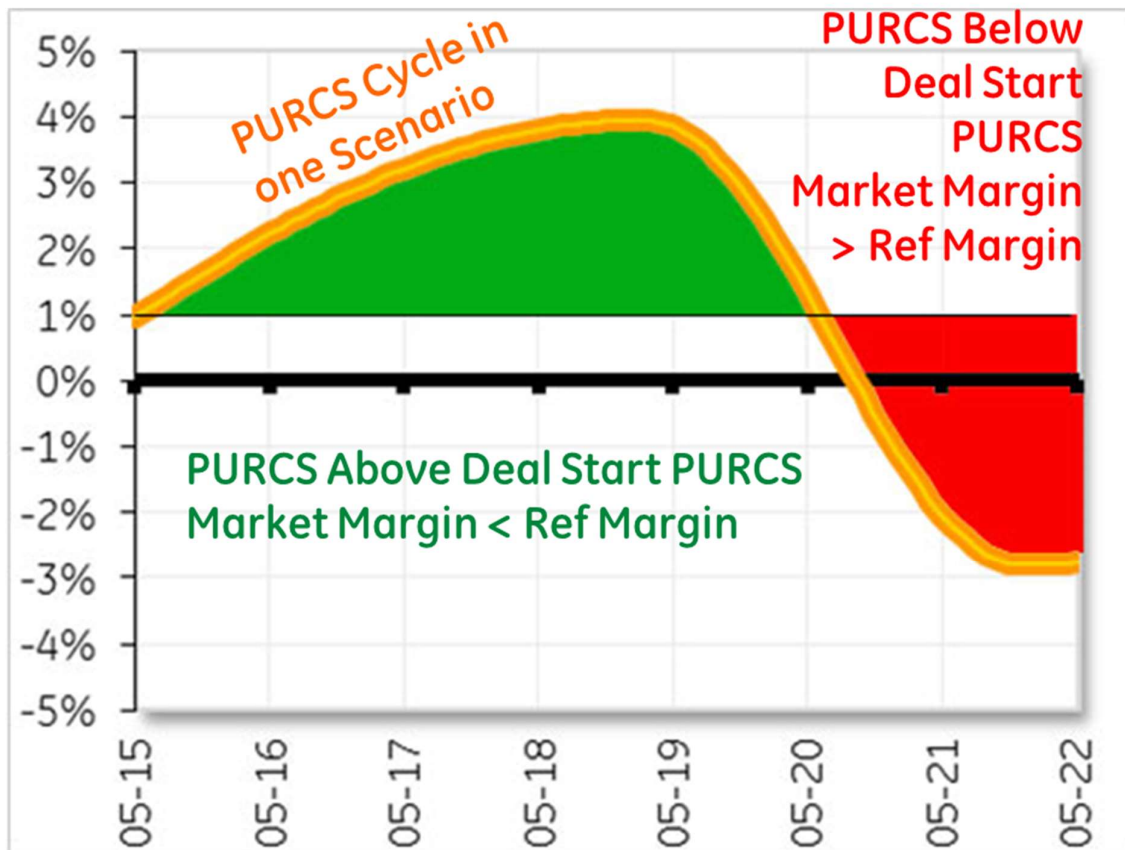
Figure 97: Yellow line shows one simulated scenario of the PURCS cycle. The green area represents the phase when the PURCS cycle amplitude is higher than at inception of the loan, with corresponding compressed market margins, creating a cost incentive to prepay, whereas the red area represents the opposite.

Cash Incentive

Again, starting from the premise that the LTV at the inception of the loan is a reference where the borrower felt that the collateral was fairly monetized given the overall terms of the loan, we look at how the LTV evolves over the tenor of the loan in each simulated scenario. We then obtain a profile over the tenor of the excess cash (or deficiency) that could be extracted in a refinancing at the same LTV as the borrower had at the inception of the loan, in each simulated scenario. The assumption that the loan can be refinanced at the same LTV is obviously a simplifying one. The aircraft will be older and market conditions may have shifted.
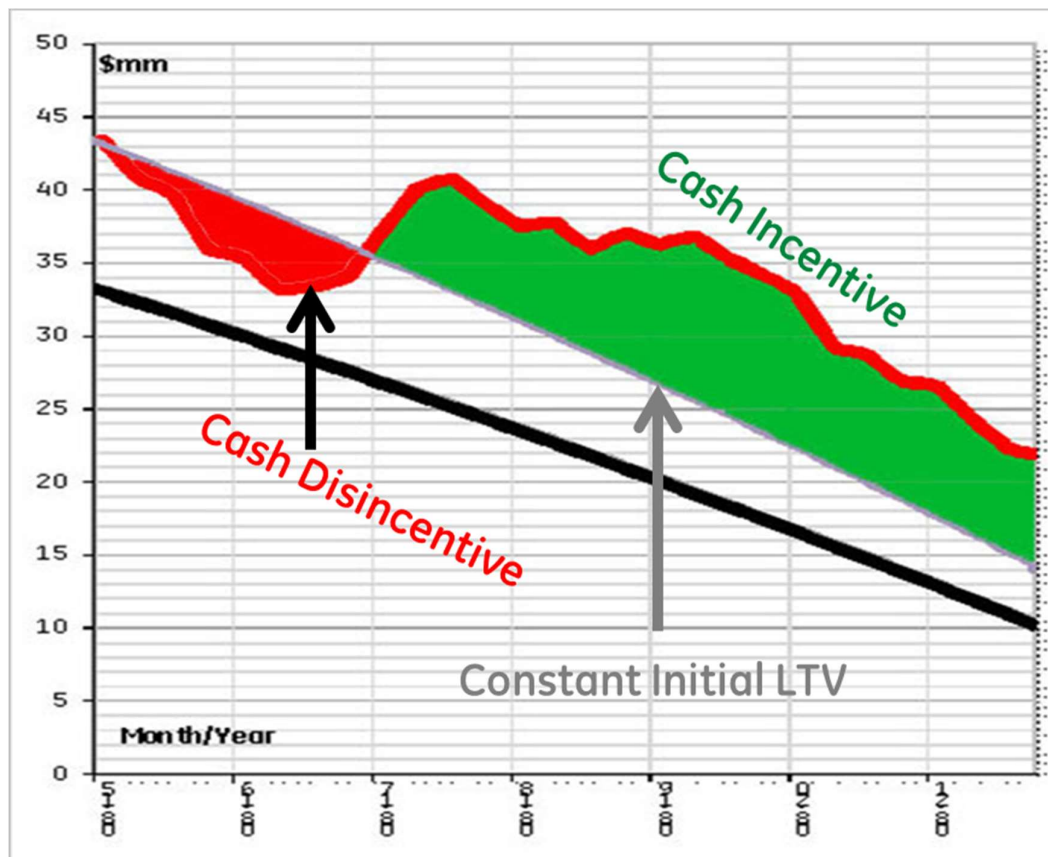
Figure 98: The black line represents the exposure profile over the tenor, the red line the simulated collateral value and the gray line the constant LTV reference collateral value. The red area creates a cash disincentive and the green area a positive incentive.

Collateral

The third "C", the collateral aspect, is not among any of the parameters that we simulate in the model. It is about the airline's fleet planning, and it is doubtful that this driver of prepayment can be predicted in a meaningful way. Of course, if the lender knows that the airline has plans to phase out a type of aircraft when granting the loan, or an investor is planning on ultimately refinancing the aircraft into a capital markets structure, we could subjectively modify the prepayment "propensity factor" (see below).

Combining the Incentives

For each simulated scenario, we now have two profiles of incentives over the tenor: one for the cost incentive expressed as dollars saved, and the other the cash inventive as dollars raised. Those are apples and oranges. So how does one weigh the two different types of incentive to obtain the overall profile of prepayment incentive? That is where the regression analysis comes in. The parameters with respect to LTV and market margin should help us to determine a "cash to cost ratio" that would permit an estimation of the combined incentive. Again, using results from the regression analysis, we can translate the combined incentive to a cumulative probability of non-prepayment for each scenario. This overall "propensity factor" has to be calibrated to empirical data and will be different for airline loans, investor loans and loans into tax enhanced transactions.

In the Monte Carlo simulation, we can now generate a uniformly distributed random variable between 0 and 1 and "shoot" horizontally into the cumulative non-prepayment probability distribution. The intercept will determine the time of prepayment. If a default has occurred

114

prior in that scenario, there is obviously no prepayment possible. Once we generate many, many scenarios, we get an overall cumulative probability distribution of prepayment over the tenor of the loan. This technique is very similar to the one the model uses for obligor default.
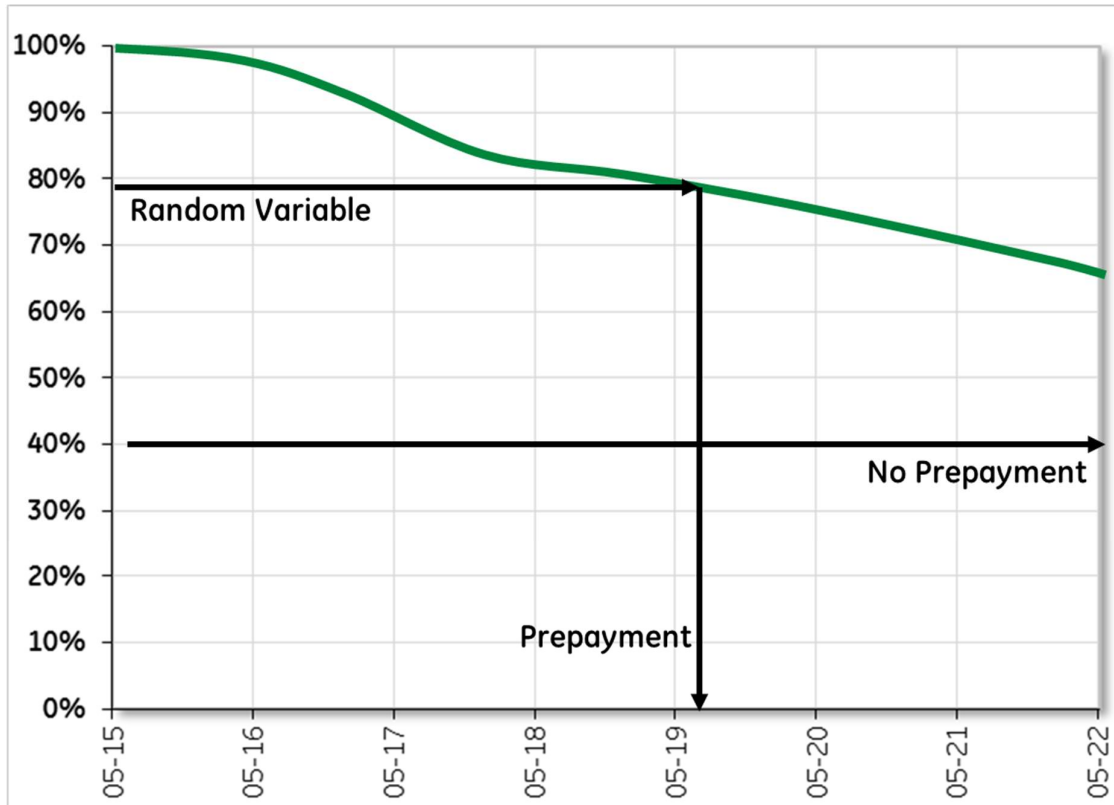

Figure 99: The green line represents the cumulative probability of non-prepayment in one scenario of the simulation. The Random Variable is "shot" at the green line to determine the time of prepayment.
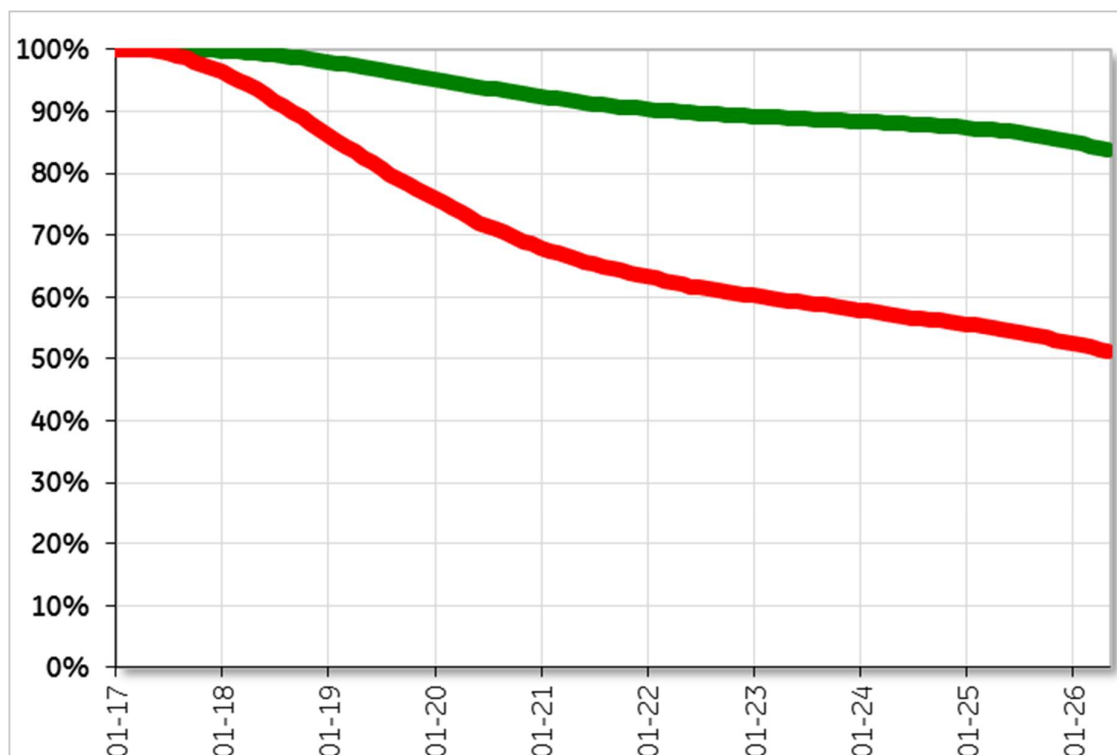
Figure 100: The cumulative probability of non-prepayment for a typical loan. The red line is a loan without any prepayment fees, the green line has a prepayment fee that is declining over the term. Note the typical shape – prepayment risk is modest initially when market conditions are close to the starting situation, then the risk increases after which the risk goes down again when the borrower has less opportunity to save cost given the shorter remaining tenor.

Valuing the Prepayment Option

Now when the prepayment events are built into the Monte Carlo model, it is very easy to measure the value of the prepayment option. By turning off the prepayment feature, one can compare the mark-to-model value with and without prepayment, and the value of the option is the difference.

Figure 101 below shows the NPV distribution of a loan run with and without the prepayment option. Note how the loss tail is the same in both cases – very rarely does a borrower prepay a loan that is in an economic loss position.

It is also easy to include prepayment terms in the optimization of the loan negotiation by varying prepayment fees or prepayment restrictions, and see how this impacts the mark-to-model value of the loan.
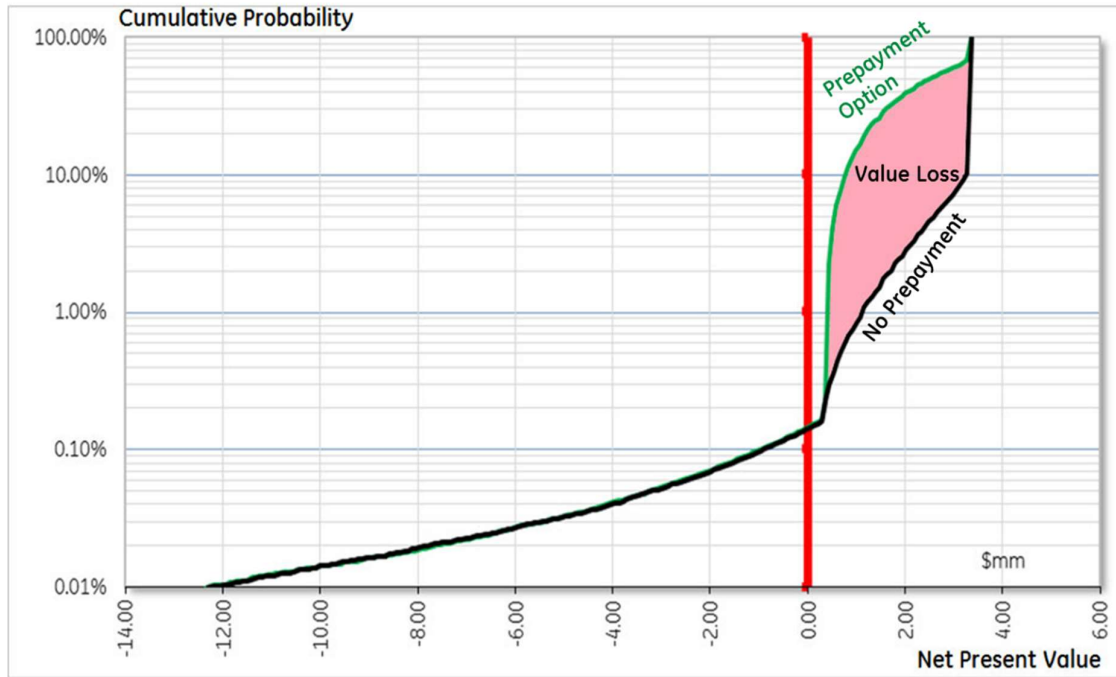
Figure 101: The NPV distribution for a loan with and without voluntary prepayment allowed. The red area represents the loss in value to the lender, hence the value of the option to prepay.

Calibration of Model Parameters

Just like with Obligor Defaults, we need to measure the discriminatory power and the level accuracy of the model. The power to discriminate between loans the will prepay and loans that will not prepay is likely to be lower than for obligor defaults. This is due to the Collateral aspect discussed above and to some degree due of irrationality in prepayment behavior (called burnout in the mortgage finance literature). The level of prepayment is strongly correlated to the PURCS cycle which is why building in strong prepayment protection for the lender is much more important when loans are granted in the trough of the cycle than at the peak. LTVs tend to be low and credit spreads high in the trough, so there is more scope for improving LTVs and compressing credit spreads driving higher cash and cost incentives for the borrower to prepay during the term.

Figure 102 below shows an example of the discriminatory power of the prepayment model based on deal count.

117
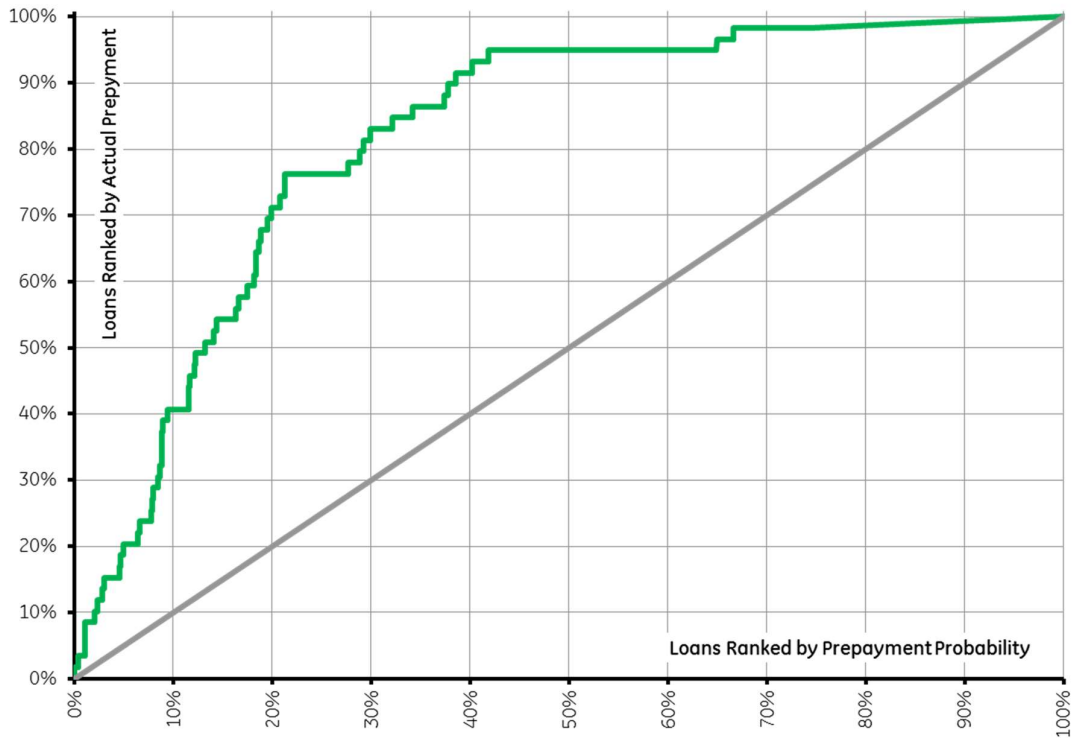
Figure 102: An example of a prepayment model's discriminatory power – in this case predicting prepayments over one quarter. This example is based on deal count, where large loans and small loans are weighted equally. The accuracy ratio (the area between the green line and the gray diagonal divided by the area of the triangle defined by the vertical axis, the horizontal 100&% line and the diagonal) is 64%.

The overall level of prepayment in a portfolio, again based on deal count, is shown in Figure 103 below.

Figure 103: Predicted versus actual prepayments in a portfolio over 6 quarters. The lines show the proportion of loans that are not p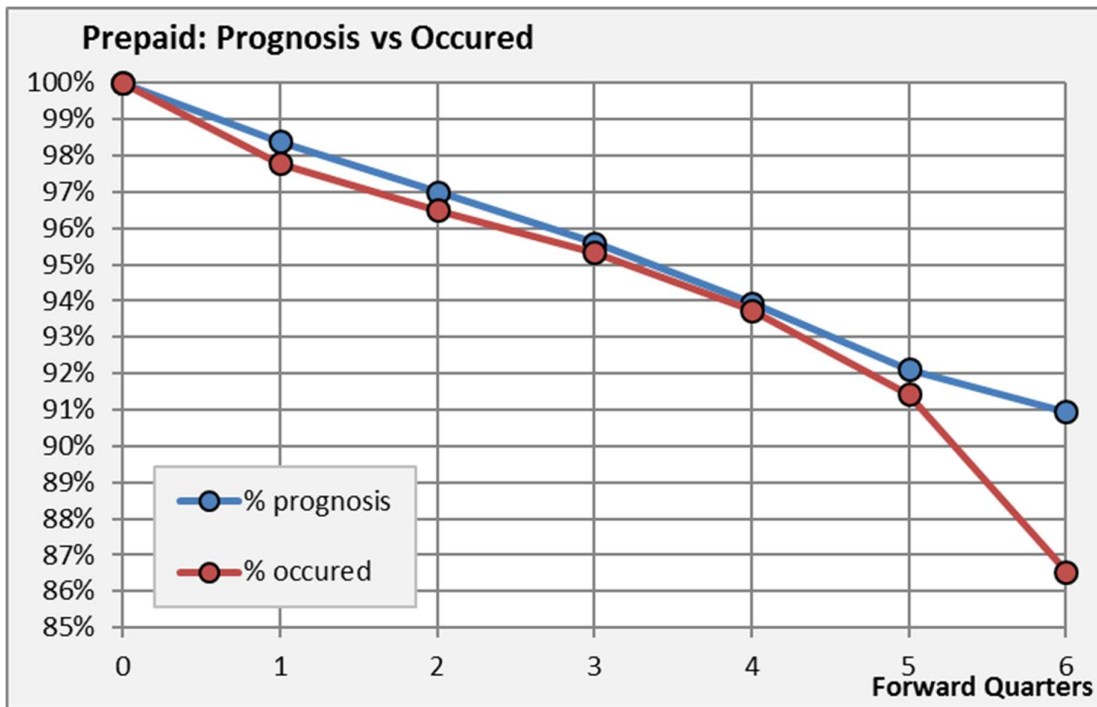repaid against predictions. Note that in this example, actual prepayments were slightly higher than predicted levels. If this persists as the data is collected, the prepayment propensity factor needs to be adjusted.

Other Risks

We have discussed credit risk, asset risk, interest rate risk, and the mitigation of these risks through diversification and cross-collateralization. Unfortunately, there are additional risks facing a lender of aircraft backed loans or lessor of aircraft.

Jurisdiction Risks

Throughout these notes, we have assumed that when the obligor makes a payment to the lender/lessor, this amount will reach its beneficiary. We have also assumed that when the obligor fails to pay, the lender/lessor can gain possession of the aircraft and benefit from any supplemental collateral. These are simplifying assumptions, and reality may not always be as straightforward. We face Jurisdiction Risk.

A country could block the transfer of funds to foreign based creditors or prevent conversion of its currency. Economic and financial conditions as a result of political actions could impact an airline's creditworthiness. Political instability or war would also impact the airline. The borrower/lessee or its country of incorporation may be involved in money laundering or terrorist financing and subject the lender/lessor to fines or loss of license to operate. Sanctions could be implemented against a country, or against specific companies or individuals in the country. The efforts of the creditor to repossess the aircraft could be frustrated by corruption, inefficiencies and bureaucracy in the court system. The de-registration and export of the aircraft may be delayed or prevented by a lack of co-operation by the various authorities, airport- or air traffic control officials. It is hoped that the Cape Town Convention (CTC) will bring a faster resolution of enforcement efforts by having expedited and firm timelines for courts to order repossession, and for aviation authorities to deregister and process export of the aircraft. This could significantly reduce the risk for the creditor and allow tighter credit margins or lower lease rentals in jurisdictions that have ratified the CTC. The jurisdiction risk can be multi-faceted: the operator base, physical location of the aircraft, and incorporation of lessee/borrower, could all be in different countries. Jurisdiction risk can to some extent be mitigated. It is crucial that the creditor conducts thorough due diligence of the counter-party and has an efficient know-your-customer process. The jurisdiction and legal system must be carefully analyzed. And the transaction and the documentation should be structured and drafted so as to protect against jurisdictional risks. Political Risk insurance is another way to mitigate against jurisdiction risk.

It is fair to say that jurisdictional risks in aircraft leasing and lending are often limited. When an airline cannot meet its financial obligations, it is often a sign that they operate too many un-profitable routes and a consensual return of an aircraft is more common than a contested repossession. A country that obstructs a legitimate repossession of an aircraft may find that its airlines will face difficulties to lease aircraft or obtain financing of favorable terms in the future. The airlines are typically an important part of a country's infrastructure.

From a modeling perspective, it is very difficult to predict specific jurisdictional risk events that may occur over the life of a loan/lease. But it is probably helpful to classify countries according to their jurisdictional riskiness and assign legal costs and downtime based on this classification. In a Monte Carlo model, this could be both expected values and a distribution. There are several providers of country scores. OECD, OFAC in the US, Transparency International and others. As the aircraft lessor/lender enjoys an easily movable collateral, these country scores do not necessarily translate to a readily applicable metric for our model. Internal data must complement any external views.

Operational Risks

Operational risks are risks arising from People, Processes or Systems. Examples are fraud, miscalculations, errors in data and assumptions, omissions in the structuring, documentation or maintenance of the transaction, or failures in systems. Luckily, the typical lender or lessor does not suffer an abundance of events of this type. So, unfortunately, there is not much empirical data to estimate or measure these risks. Operational risks are often related to credit risks, or asset risks; maybe there was a defect in the security package or local counsel forgot to file that mortgage, but the borrower did not default, so the problem never surfaced!

Liquidity

Liquidity is a word that has been used a few times since the Lehman bankruptcy in 2008. What does it mean and what are the implications for aircraft loans and leases? I think there are two aspects to liquidity that are relevant for us:

1) **Asset Liquidity**: The liquidity of the asset (loan or lease) itself, how quickly does it trade and what transaction costs are involved?

2) **Funding Cost**: What is the cost of funds for the holder of the asset expressed as a premium over the benchmark (libor or swap), and how is that relevant to the pricing or valuation of the loan or lease?

Clearly, an investor would seek a higher return from an asset that has high search costs, high transaction costs, and a longer time lag required to complete a trade. This aspect has been ignored in the analysis in this paper. If the Mark-to-Model value is derived from a "market" Risk-Reward Ratio, we have to be careful and make sure the reference market is based on assets with similar liquidity. In the aircraft finance market, one could compare the pricing of EETCs or ABSs to the ones of private bank loans or outright sales of aircraft with leases attached. We would have to adjust for the diversification and cross-collateralization that is embedded in these structured securities. This is another field of further research!

The funding cost was a non-issue for many banks and large finance companies before the "financial crisis". Those institutions typically paid the reference rate, and it was a correct valuation basis to use this reference rate. Any margin above was income. However, when the crisis hit, funding costs shot up, and the margin over the reference rate was no longer an adequate measure of earnings and a basis for NPV. It could be argued that any spread over the risk-free rate that the asset holder pays is unrelated to the value of the assets, and just a reflection of the fact that the asset holder is paying a premium over the reference rate for the right to go bankrupt. It would make absolutely no sense for a lender to grant a loan to a borrower with a loan margin that is below, or even near its funding cost (spread over the reference rate). The loan would have no value for the lender. Yes, if there were buyers of the loan that had lower funding cost, the loan could be sold, but again, transaction costs would be high. In a world where most financial institutions pay a premium for its funding, over and above the reference rate, a more correct benchmark for valuing loans and leases would be the ***marginal funding cost facing prime players in this market***. Since the spread over the benchmark eats into the lender's or lessor's margin to their customer, the risk-reward profile has to be adapted to the realities of the costs of funding. Only very financially strong creditors, and hence with access to inexpensive liquidity can operate in the low part of the spread spectrum, the weaker ones will have to migrate to the higher spreads, with corresponding higher risks.

Lender's Liability

Lender's Liability can arise in several ways. It can be thought of in terms of the exposure to claims that might be taken against a lender. This might arise from the lending relationship or process itself (for example alleged misrepresentation or interest rate manipulation) or from being in possession of the asset after a default (mortgagee in possession) or perhaps from taking control of the borrower's business by enforcing security in circumstances where the lender effectively decides to run the borrower's business.