Stochastic Discount Factors and Real Options

Paul Edge

May 20, 2011

Abstract

This paper uses the stochastic discount factor (SDF) to price real options and introduces the expected discounted shortfall (EDS) risk measure to control risk.

A multivariate covariance based SDF modelling framework is described. Explicit formulae linking the correlation matrix to the risk premium are derived for assets prices following both Brownian and Ornstein-Uhlenbeck processes. Applying the SDF to real option problems simplifies calculations and economic assumptions by removing the requirement for replicating portfolios.The SDF method does not identify a hedging portfolio, so other risk control methods have to be used to compensate. EDS is a coherent, multi-period risk measure that calculates the present value of the risk to the shareholder, that cashflows are insufficient. An example real option is included, focusing on cashflow, using the SDF approach to measure the change in return and the EDS risk measure to price the increase in risk. The changing risk and return profile over time is also studied.

Keywords: Stochastic Discount Factor, SDF, Real Option, Expected Discounted Shortfall, EDS, Risk Measure, Coherence

Contents

List of Tables

*Nomenclature

$\underline{C}$ Cashflow space

CE^i Certainty equivalent adjustment for asset i

$Cov\,(X,Y)$ $= E\,[XY] - E\,[X]\,E\,[Y] =$ Covariance between X and Y

$\mathbf{M}$ SDF vector

$E\,[\,]$ Expectation function

$\mathbf{I}$ Identity matrix

$\mathbf{P}$ Matrix of prices

$\Pr$ Probability vector

$price\,(\,)$ Pricing function

$\Pr\,(x)$ Probability of x occurring

R^f Risk free rate

R^i Return on asset i

${}^eR^i$ $= \frac{R^i}{R^f} =$ Excess return over the risk free rate of asset i

R^p Return of the market

$R_{s:t}$ Return between time s and t

$E^P\,[\,]$ Expectation taken using the **real world** probability measure

$E^Q\,[\,]$ Expectation taken using the **risk neutral** probability measure

s Start time

m Stochastic discount factor

t End time

$V\,[X]$ $= E\left[X^2\right] - E\,[X]^2 =$ Variance of X

$\mathbf{X}$ Matrix of cashflows

x Cashflow

β^i Market beta for asset i

β_t Intrinsic discount process at time t

ρ Correlation coefficient

α Mean reversion parameter

ω Possible state of the economy

$\sigma\,(\,)$	Risk measure (function)
Ψ	State price
ε	Random element
θ	Vector of portfolio weightings
σ	Volatility parameter (variable)
i	Asset index
T	Time index
CDP_t	Cashflow default point. The minimum outgoing cashflow at time t
ROA	Real options analysis

1 Introduction

An option is the right, but not the obligation to take action at pre-determined cost for a pre-determined period of time depending on how uncertain conditions evolve. A real option is any investment in a physical or intangible asset, in human competence, or in any capabilities that help an organization to envision and respond to future contingent events.

Valuing an asset using (possibly mean reverting) random processes to generate future prices is a widely accepted approach to asset valuation, but the choice of a suitable discount rate is a contentious issue. This paper considers the discount rate as a random process, called a stochastic discount factor (SDF). A simple but generic approach to SDF modelling is introduced, based around correlations. This framework is then applied to real options pricing and to risk calculations. The Monte Carlo SDF approach of this paper values real options and removes the need for an explicit replicating portfolio, simplifying calculations and economic justifications.

The risk of holding the real option through time must still be managed, so this paper introduces the expected discounted shortfall (EDS) measure which is a coherent, multi-period risk measure that is presented as a cost from the shareholder's point of view.

1.1 Stochastic discount factors

The stochastic discount factor (SDF) is a stochastic process that discounts a projected future cashflow to give a present value. The SDF process takes a different value at each point in time and is also dependent on the state of the economy at that time. Only one SDF process is applied to all cashflows, pricing all assets simultaneously. The SDF is synonymously referred to across the academic literature as the state price deflator[1], stochastic discount factor, state price density, state price kernel and the marginal rate of substitution.

Looking at a project's cashflows in isolation, the SDF method can calculate how much value taking on a project would add to a diversified investor. Each opportunity to influence the project's cashflows is a real option, and a chosen combination of real options is a strategy. SDFs allow the value of different strategies to be compared and decisions to be ranked. Although the SDF can value a project, it does not directly quantify the risk of investing in a project.

This paper uses covariance to simplify the mathematical complexities of other SDF implementations, formulating a generic SDF model that explicitly links the correlation matrix to risk premiums via the SDF. The approach is simple to understand and implement alongside existing models, and can be achieved without sacrificing any descriptive power. The approach allows many different types of stochastic processes to be combined without increasing the volatility of the SDF.

The derivation of a log-normally distributed, mean reverting, model uses the correlation matrix to fully describe the relationship between the risk premium of each asset and the SDF. The model in this paper is simple and tractable, allowing results to be checked easily.

1.2 Real options

A real option valuation method is needed that can price all types of assets, uses the real world probability measure, and is simple for practitioners to use.

This paper introduces a new approach to valuing real options, using the SDF within a Monte Carlo framework. An equilibrium based theoretical argument is used to justify the approach, avoiding many assumptions that, to date, have hampered the general acceptance of real options.

[1] Not to be confused with the economic, GDP deflator which is a historic measure of changes in prices in an economy.

SDFs can be used to disguise the mathematical complexities, allowing a practitioner to value real options without knowledge of the techniques needed to generate the stochastic variables. SDFs also remove the complexity and subjectivity around choosing a discount rate.

Under the SDF approach there is only one category of real option model. For example, there is no difference between implementing a 'switching' option model or an 'abandonment' option model. The underlying SDF model and implementation method stays the same whatever option is valued, only the cashflows change. SDFs encourage real options to be viewed and valued collectively, allowing overall strategies to be measured against one another, on a combined risk and return basis.

Combining real options, the SDF, and EDS allows the risk and return to a company of a given real option strategy to be fully defined. Real options calculate the return from a project and the EDS calculates the risk to the firm when a project is commissioned.

1.3 Risk pricing

The Monte Carlo SDF approach does not produce a hedging portfolio, so the risk of holding a real option must be measured using other methods. If a company had an infinite ability to raise cash then risk would be ignored and all valuations could be done on an expected value basis. As the cash limit is set at the company level then this is where the risk (of a real options strategy) should be measured and controlled.

This paper defines a new, coherent, multiperiod, value based risk measure that is applied at the company level. EDS is the only risk measure that is able to coherently combine risks from multiple time periods into a single present value. It produces an absolute price for a risk which can be subtracted from the asset value. Mapping the risk dimension onto value makes EDS the ideal measure to judge the effectiveness of different strategies from a shareholder's point of view.

1.4 Paper structure

The literature on real options, SDFs and risk measures is summarised in section 2.

Section 3 documents the SDF modelling approach, from first principles through to a generic covariance SDF model. The end of the section describes an analytic SDF model where the random element of the return is normally distributed and is controlled by the correlation matrix.

Section 4 discusses real options and the important role that SDFs play in simplifying theoretical justifications and computations by removing the requirement for a replicating portfolio to be found. The section discusses how equilibrium arguments are relied upon for theoretical justification.

Section 5 discusses how the risk of a particular real option strategy can be measures and introduces the EDS risk measure.

Section 6 demonstrates the covariance SDF model from section 3 and why the EDS risk measure from section 5 should be included in real options analysis.

2 Literature Review

Section 2.1 collects the various literature surrounding SDFs, comparing approaches and closes by detailing this paper contributes. Section 2.2 looks at risk measures. Section 2.3 critiques the common theoretical methods used to value real options. Section 2.4 highlights the gaps that this paper attempts to fill.

2.1 Stochastic discount factors

A contingent claim asset is an asset that delivers a payoff if a certain set of conditions occurs. State assets are a particular type of contingent claim asset, introduced by [Arrow & Debreu, 1954], that pay a unit of currency if a particular state of the economy occurs, and nothing otherwise. [Varian, 1987] gives a good introduction to state prices.[Etheridge, 2002] explains why weak arbitrage is impossible if and only if there exists non-negative state prices. [Girotto & Ortu, 1997] provide a mathematically rigorous description of the equivalence between SDFs and state prices in discrete finite space and time, also showing that there is a single SDF if and only if a market is complete. In continuous time [Harrison & Pliska, 1981] showed that absence of arbitrage is equivalent to the existence of an intrinsic discount process that makes the discounted asset process into a martingale. [Jarvis et al., 2001] provide an excellent introduction to SDFs. [Cochrane, 2004] provides the links between some standard asset pricing models, showing that many popular asset pricing models are special cases of the SDF model.Thorough SDF literature surveys can be found in [Ferson, 1995], [Smith & Wickens, 2002] and [Campbell, 2003].

[Chapman, 1997] approximates the SDF using orthogonal Legendre polynomials, creating a continuous function of a set of observable variables. [Rosenberg & Engle, 2002] use Chebyshev polynomials and allow the shape of the SDF to be flexible over time. Additional state variables can be added using tensor products. This approach is flexible and able to price non-linear assets, like derivatives. More than one asset can be priced simultaneously but as the dimension size increases the tensor products quickly become burdensome.

[Rubinstein, 1976] derives a simple form of the SDF. [Epstein & Zin, 1991] use the reciprocal of the return on the optimal portfolio which is hard to calculate when there is a large number of assets. [Arajo et al., 2005] apply panel-data techniques to asset returns to construct a consistent estimator for the SDF. [Ait-Sahalia & Lo, 1998] use equity index option prices to non-parametrically estimate the average SDF, which is used in [At-Sahalia & Lo, 2000] to give a VaR distribution adjusted for risk aversion. [Breymann et al., 2005] show that diversified indices such as the FTSE All Share, S&P 500 or Morgan Stanley Capital Growth World Index (MSCI) are all appropriate market proxies for the theoretical 'growth optimized portfolio' introduced by [Platen, 2006], which is equivalent to the SDF.

[Hansen & Jagannathan, 1991] construct bounds for SDF from historic times series, with and without a risk free asset. [Hansen & Jagannathan, 1997] construct a benchmark linear SDF model against which other models can be tested, giving a measure of how well prices are calculated by the model. This method is used in [Alvarez & Jermann, 2005] to put empirical bounds around each part of the SDF.Unfortunately the non-parametric approach only measures the SDF for the particular underlying asset that provided the data. The results will not necessarily generate an SDF that can be used to price other assets.

[Duffie et al., 2000] define an affine jump diffusion state process for pricing defaultable, fixed income bonds and options with stochastic volatility. This structure of the SDF allows parameter estimation via Laplace and Fourier transforms, both of which can be used to give closed form solutions to option prices. The SDF process is actually a linear factor model of underlying processes, not a process in its own right. There are various affine models in the literature, for example [Bekaert & Grenadier, 2002] look at equity pricing and bond pricing, [Dai & Singleton, 2000] study the term structure, [Piazzesi, 2009] surveys affine bond models and [Monfort & Gourieroux, 2007] give a model for stochastic interest rates. Generating the SDF as weighted portfolio of other assets is flexible, but the correlation between each asset and the SDF, and therefore the risk premium of each asset, is hard to control.

[Constantinides, 1992] was the first to define term structure between numeraries using state prices and [Rogers, 1997] extended the approach to SDFs. [Smith & Speed, 1998] model all processes as SDFs from the same general equation creating a multi-currency, continuous time and space, full yield curve model. The price of an asset is defined by the ratio of two SDFs. [Platen, 2000] takes a similar approach but with greater mathematical rigor. [Bakshi et al., 2008] compare SDFs in three countries, fitting empirical option data to parameters of a foreign exchange model with a stochastic time change diffusion component for global risk and an independent stochastic time change jump diffusion component for country specific risk. The exchange rate is the ratio of two SDFs. [Mo & Wu, 2007] take a similar approach except that they use data from equity option markets and place the jump risk in the global component of the equity model instead of the country component. Defining orthogonal global and independent country components means that the only co-movement between countries is through the global component. This does not scale well to multiple economies where there could be dependencies between countries that are unrelated to the global risk. It is hard to generate correlations in multivariate stochastic time change models when using stochastic time counters.

Modelling all prices as a ratio of two SDFs is a complex way to approach asset pricing. It is hard to control and calibrate ratios of two stochastic processes and calibration to the market price of derivatives written on those assets is even harder. Using the ratio modelling approach means that adding extra assets increases the SDF volatility, increasing the convergence time when implemented using the Monte Carlo method. An alternative approach to multi-currency modelling is given by [Hadzi-Vaskov & Kool, 2008] who derive the SDF in foreign currencies by modifying the local currency SDF using exchange rates retaining control over the exchange rate process and keeping the nature of the original SDF model.

A good deal is a trade that any investor will be willing to make, either because risk to return ratios are large or arbitrage opportunities exist. The good deal bound approach puts an upper and lower bound on values for the SDF. These two sets of values are then used to give an upper and lower price for a real option. Good deal bounds are not equivalent to pricing with pure Sharpe ratio arguments as they can bring extra market information into the equation, such as correlation with a traded asset. [Amram & Kulatilaka, 1999] suggest that real option practitioners should try to find a portfolio of traded securities that is correlated (even if just partially) with the untraded risk. Pricing via correlation with similar assets is a common approach when trading financial options on illiquid assets. [Staum, 2004] gives a rigorous description of how SDFs can be used with good deal bounds, finding minimum and maximum values for options based on realistic values for their Sharpe ratio, after taking into account correlations with other assets. [Bernardo & Ledoit, 2000] use SDFs to calculate good deal bounds for options, and [Cochrane & Sa-Requejo, 2000] apply the same approach to real options, deriving an analytical solution similar to Black Scholes but with extra terms. [Chapman et al., 2001] use SDFs to value cashflows to all stakeholders in a pension scheme identifying the winning and losing parties. The good deal bound approach to real options has also been adopted by [Thijssen, 2002] in which he compares the utility based approach and the good deal bound approach.

Disadvantages of current SDF approaches Non-linear SDFs are able to price complex assets, like derivatives, but are unsuitable for multiple sources of risk. Employing stochastic volatility with jump processes to produce non-normal returns does not allow correlation between assets to be controlled. The exponential affine jump diffusion approach overcomes the above problems but only allows Gaussian processes to be combined with jump process.

2.2 Risk measures

Definition A risk measure assigns a real number to a random variable.

[Artzner et al., 1999] define the properties of a coherent risk measure. Coherence is essential because it ensures that a risk measure produces consistent decisions. If a risk measure is not coherent then is impossible to know if a result is optimal or not. For example, conflicting results could occur at different confidence levels, or further diversification of a portfolio may lead to an increase in the risk measure. [Ng & Varnell, 2003] define a coherent risk measure that combines with SDFs to value the effect of frictional costs. They recognise the difference in pricing between systematic and unsystematic risk and show that frictional costs can help explain why some diversifiable risks attract a risk premium. [Csoka et al., 2007] extend the idea to show that *any* risk measure using SDFs is a coherent risk measure, for all cashflows. This property is used in section 5.

[Froot & Stein, 1998] say that hedgeable risks should not need any risk capital, because they can be mitigated in the market. Unhedgeable risks can be managed by holding an amount of risk capital that can be attributed to the risk. The amount of risk capital needed for the unhedgeable part of the risk is dependent on what other unhedgeable risks are held in the portfolio. The same approaches that are used to handle the unhedgeable risks in a trading operation can be used to handle the risks of real options. That is, the real option risks can be combined, measured and compared at the company level. [Ho & Liu, 2002] value the government guarantee of debt as a real option of a project. The risk of different strategies can be evaluated by noting the change in value of this real option.

[Ye & Tiong, 2000] presented the NPV@Risk method to value privately financed infrastructure projects. One of the main advantages of this approach as a risk measure is that risks in different time periods can be combined. The NPV@Risk approach takes the risks from all time periods and discounts them to a present value, solving the time period problem. Unfortunately, the rate at which cashflows are discounted is not defined, nor is NPV@Risk a coherent risk measure.

2.2.1 Expected shortfall (ES)

EDS, introduced in this paper, is an extension of expected shortfall (ES), also referred to as conditional value at risk which was introduced by [Acerbi & Tasche, 2002]. It is defined as the expected value of the cashflows exceeding VaR_γ so

$$ES_\gamma(X) = \frac{1}{1-\gamma}\int_{X \geq VaR_\gamma} xPr(x)\,dx = VaR_\gamma + E\left[X - VaR_\gamma | X > VaR_\gamma\right]$$

where X is the random variable being measured and $Pr(x)$ is the probability that $X = x$. Expected shortfall satisfies all the properties of a coherent risk measure, but the measure is not cumulative, so cannot be used for multiperiod risk calculations.

.

2.3 Real options

When making financial decisions there is a need to convert uncertain future cash flows into a present-day value. Many methods have been developed to achieve this.

[Sick, 1995] uses the [Miller & Upton, 1985] principle which assumes that an amount of exhaustible natural resource is a function only of its spot current price, net of extraction cost. Ignoring discount rates may provide a quick valuation, but ignoring the time dimension is not a credible approach to option analysis. Deciding when to exercise an option is a fundamental part of flexibility.

The idea of discounting expected cash flows at a rate adjusted for risk was proposed in [Fisher, 1907]. The Net Present Value (NPV) concept was initially developed to value bonds and stocks by passive investors who required a different risk adjusted rate for each asset. It is common for all risks factors to be combined into a single, suitably weighted, risk adjusted, discount rate, often set with reference to the capital asset pricing model (CAPM) (see [Sharpe, 1964]) where the risk premium on the asset depends on its correlation with the rest of the market.

The NPV approach is powerful, simple and well understood but based on value maximization without taking into account uncertainty and flexibility within an investment. All risks are discounted at the same (possibly time dependent) rate, whether the risk is correlated with the market or not. It makes assumptions about future scenarios, but does not allow different responses to be made in different scenarios. The flexibility within a project can be a major source of value.

Using a single, constant, discount rate assumes that the riskiness of cashflows increases at a geometric rate over time. This is fine for projects that are time homogeneous and memoryless, but this is generally not the case (see [Trigeorgis, 1996]) so projects require different rates over different time periods.

Attempts to modify the NPV approach to account for flexibility are usually unsuccessful. [Myers, 1976] states that ”If NPV is calculated using an appropriate discount rate, any further adjustment for risk is double counting”. Unfortunately, using the NPV approach with further adjustments for risks is common practice. [Mathews et al., 2007] use a Monte Carlo projection of prices with a fixed discount rate, but allow different types of cashflow to be discounted at different rates. [Swinand et al., 2005] also make adjustments to the NPV approach, saying that ”an important insight of real option theory is that the cash flows are appropriately discounted at a rate based on the weighted average cost of capital, but the value of flexibility should be discounted at the risk-free rate of interest.”. Decisions based upon these ’quick and dirty’ methods are ill founded. [Hodder et al., 2001] compare the risk adjusted, risk neutral and certainty equivalent approaches. They show that risk adjusted discount rates can be made to work but only retrospectively.

The market asset disclaimer (MAD) approach developed by [Copeland & Antikarov, 2001] treats the project and real option as if they were actually traded. The best replicating asset for the option is assumed to be the project without flexibility, which is valued using the NPV method. The need to find an actual traded replicating portfolio is removed because Copeland & Antikarov reason that no ”asset has a better correlation with the project, than the project itself.” The NPV approach does not value risks properly, so correctly pricing the project without flexibility (e.g using the correct discount rate) is just as complex and important as valuing the flexibility. No attempt is made under MAD to control risk when executing the real option. The approach requires decision trees to be built to value the options, so is not appropriate for complex, multivariate, compound options. [Borison, 2005] points out that, under the MAD approach, the value of the investment underlying the real option is based entirely on subjective data. It is only tied to market data by the discount rate and ignores the possibility that an actual replicating portfolio exists.

A popular valuation approach is to apply risk neutral probabilities to price real options. The assumption of risk neutrality, and the use of the risk free discount rate, is justified by identifying a portfolio of traded investments replicating the cashflows of the real option, removing all risks from the asset (also assuming perfect and frictionless markets) making the asset equally valuable to all investors. Using no arbitrage arguments (see appendix A.2), the value of the option is the value of the replicating portfolio. This approach to real options works well where underlying risks are freely traded (e.g. with commodity based real options), but the risk neutral approach is harder to justify when the risks cannot be hedged in the market. Many assets are not freely traded and there are few very long term traded options available to construct hedges. [Dixit & Pindyck, 1994] only use the risk neutral approach when investment values can be reasonably tracked by publicly traded assets. Decision analysis is used when investments are dominated by corporate specific or untradeable risks. Each real option type is priced using a different model.

Utility functions are used to define risk preferences of investors, which are then used to justify a risk premium on a real option. Utility functions were introduced to real options by [Henderson & Hobson, 2002], followed up in [Henderson, 2007] and with mean reversion in [Yang & Ewald, 2007]. They often aid the finding of analytic solutions to problems. [Miao & Wang, 2005] use a utility function to establish certainty equivalent values. Section 4.1.3 demonstrates why utility functions should not be applied to decision making within publicly traded companies.

The concept of certainty equivalent values was introduced by [Theil, 1957] and was ‘popularized by [Robichek & Myers, 1965]. An adjustment is made to an expected cashflow, so that it can be discounted at the risk free rate. The method projects prices using real world probabilities, but each cashflow requires its own correlation with the market portfolio to be estimated. [Copeland & Antikarov, 2001] recommend certainty equivalent values for real option valuation as it complements their MAD approach. They show that the certainty equivalent approach is the same as the risk adjusted approach.

At [University of Maryland, 2002] Eduardo Schwartz said ”I don’t see companies using real options to value projects where the risk-neutral distributions have to be obtained using an equilibrium model.”. For a company to accept the equilibrium approach, real world distributions should be used. Realism requires that models are able to incorporate drivers of risk that are likely to be correlated with the economy (such as uncertainty about price, quantity, and operating costs) and be able to account for other risks which are unlikely to be correlated with the economy (like technological uncertainty). As much market information as possible should be involved, as the analysis would be incomplete if information contained in market prices was ignored. Proper treatment of risk from the shareholder’s perspective is one of the key features required of real options. [Borison, 2005] categorises real option approaches into four distinct groups; whether probabilities are derived from the market or made by subjective assessment, and whether risks are publicly traded or not. [Mello & Pyo, 2002] discount untraded risks at a risk adjusted rate, not the risk free rate. Probability assessments of untraded risks are made subjectively by the managers of the firm (but the managers’ risk tolerances are ignored).

Many prices display mean reversion, which [Metcalf & Hassett, 1995] apply to real options. [Sarkar, 2003] adds correlation with the market (although no beta calculation is given) and [Jaimungal et al., 2009] apply mean reversion to both investment amount and project value, driven by Brownian motions.

SDFs can help in real options analysis by simplifying the modelling method, allowing complex situations to be analysed and results to be explained.

2.4 Gaps in the literature

Creating replicating portfolios is impossible to do with untraded assets. Many real option approaches are rejected because they require a replicating portfolio to be defined. The certainty equivalent approach does not rely on replicating portfolios, but is impractical for complex options because each cashflow requires its own covariance with the market to be calculated.

If a replicating portfolio does exist then the risk neutral approach can be applied to real options, but this approach assumes that all asset growth is at the risk free rate. Real world probabilities are needed so that realistic scenarios are projected, and because the parameters of non-normal distributions change when translated from the real world to the risk neutral measure. The SDF method uses real world probabilities, generating realistic growth in prices that is easier to explain than risk neutral projections.

The good deal bound SDF approach to real options (literature review section 2.1) allows upper and lower analytical bounds to be placed on the value of a real option, but for complex, multidimensional problems (when solution speed is not vital) only the Monte Carlo SDF approach of this paper is practical.

Risk measures based on percentile estimates are only coherent when normal distributions are used. Other coherent risk measures, such as expected shortfall, are defined at a single timepoint so they cannot measure risks over multiple time periods. No risk measure is defined with reference to an absolute cash level that also takes into account the correlation of a risk with the market. SDFs can be used to create a coherent risk measure that spans multiple time periods and takes into account correlation with the market. The current methodology is seen by practitioners as being too complex. A simpler approach is required.

3 Stochastic discount factors

Introduction Throughout this paper the process m will be referred to as the stochastic discount factor (SDF). Across the academic literature it is synonymously referred to as the state price deflator[2], stochastic discount factor, state price density, state price kernel and the marginal rate of substitution. This section blends material from [Cochrane, 2004], [Duffie, 2001] and [Jarvis et al., 2001].

Section 3.1 discusses a special type of contingent claim called a state asset. In particular it shows how prices are calculated when time and state space are discrete, where the concepts introduced are easier to understand. When there are more states than assets the market is said to be incomplete. Section 3.2 introduces probabilities and the concept of the SDF, defining prices in incomplete markets. Section 3.3 shows why the SDF is arbitrage free and extends the concept to continuous time. Section 3.4 shows that risk and return are controlled by covariance with the SDF.

The final section 3.5 contains novel contributions to the SDF literature. It introduces a simple way to build a generic, multivariate, covariance based SDF model which is easy to understand and quick to implement within a Monte Carlo framework. The section ends with two example SDF models; a log-normal, SDF based pricing model and a similar model that adds mean reversion.

3.1 State prices

3.1.1 Contingent claims

A contingent claim asset is an asset that delivers a payoff if a certain set of conditions occurs. It pays a predefined amount in the future if the world enters a certain state, where a state can represent any collection of events occurring. For example, the payoff claimed by the holder of a put option on a share is contingent on the share being below a certain value. As contingent claims can cover any event, it is possible to redefine all assets in terms of contingent claims.

3.1.2 State assets

State assets are a particular type of contingent claim asset that agree to pay 1 unit of currency if a particular state occurs, and 0 if any other states occur.

Define $\Psi(\omega)$ to be the current price of a state asset that pays 1 only if the future state ω occurs. $\Psi(\omega)$ is said to be the state price of a ω state asset, or more generally, Ψ is referred to as the state price. If state ω occurs in the future then the value of the state asset at that time will be $\Psi(\omega) = 1$, if any other state occurs then the state asset will be worth 0. These state assets can be used as building blocks to create a portfolio of state assets so that the payoff of the portfolio in all future states exactly matches the payoff from any contingent claim asset.

Let $x(\omega)$ be the future value of a contingent claim asset x if the world ends up in the state ω. The state price equation can be defined as

$$price(x) = \sum_{\omega} \Psi(\omega)\, x(\omega) \tag{1}$$

where $price(x)$ is the market price now of the contingent claim asset x. The current price of this portfolio of state assets should match the current price of the chosen contingent claim asset. If the price differs then an arbitrage opportunity will exist (see appendix A.2 for a definition of arbitrage).

[2] Not to be confused with the economic, GDP deflator which is a historic measure of changes in prices in an economy.

3.1.3 Risk free assets and state prices

An asset is risk free if its payoff is independent of the future state, so $x(\omega) = Constant$ for all ω.

Consider the amount paid now in exchange for receiving a guaranteed amount of $x(\omega) = 1$ in the future, so $price(1) = \sum_{\omega} \Psi(\omega) 1$. Enforcing non-negative interest rates requires that $\frac{1}{R^f} \leq 1$ so

$$R^f = \frac{1}{price(1)} = \frac{1}{\sum_{\omega} \Psi(\omega)} \geq 1 \Rightarrow \sum_{\omega} \Psi(\omega) \leq 1. \tag{2}$$

To avoid arbitrage all state prices must be non negative, so $\Psi(\omega) \geq 0$ for all ω. If a state price were less than 0 then it would mean that someone would pay you to hold an asset that will either expire worthless, or pay 1.

3.1.4 State prices in discrete time and space

When time and space are discrete and finite, uncertainty can be defined as a finite set $\omega \in \{1, \ldots, S\}$ of S future states, one of which becomes true. Assume there is an economy that has N assets and their cashflows in each state can be represent by an $N \times S$ cashflow matrix $\mathbf{X}$ where the element $x_{i\omega} = x_i(\omega)$ represents the amount paid by asset i in state ω. Define the prices of the N assets as the vector $\mathbf{P}$. This means that a portfolio vector θ has market value $\mathbf{P^T}\theta$ and payoff vector $\mathbf{X^T}\theta$. A vector of state prices $\mathbf{\Psi^T} = [\Psi_1, \ldots, \Psi_S]$ is a weighting such that the price of an asset can be equated to its payoffs, so $P_i = \sum_{\omega=1}^{S} x_{i\omega}\Psi_\omega$ for all assets i. It is possible to think of $\Psi_\omega = \Psi(\omega)$ as the cost of obtaining an additional unit of currency in state ω. Extending the equation to the vector of prices $\mathbf{P}$ gives the matrix representation $\mathbf{P} = \mathbf{X\Psi}$.

If there are as many assets as states then the state price equation $\mathbf{P} = \mathbf{X\Psi}$ can be reversed by inverting the cashflow matrix, $\mathbf{X}^{-1}\mathbf{P} = \mathbf{\Psi}$ uniquely defining the state price vector $\mathbf{\Psi}$. This is equivalent to solving a set of simultaneous equations, so for a unique solution the matrix must be invertible, i.e. the payoffs are linearly independent.

For a payoff cashflow matrix $\mathbf{X}$, price vector $\mathbf{P}$ and portfolio vector θ, the combination is weakly arbitrage free if $\mathbf{P^T}\theta \geq 0$ whenever $\mathbf{X}^T\theta \geq 0$. Farkas's Lemma implies that a system is weakly arbitrage free if and only if there exists non-negative state prices $\mathbf{\Psi}$ such that $\mathbf{P} = \mathbf{X\Psi}$.

3.1.5 Numerical example

What follows is a simple example showing how state prices are constructed. The example economy has two future states and two assets with known current market price and future cashflows.

Price/Payoff matrix	Current Price	Boom state payoff	Bust state payoff
Asset A	1.65	3	1
Asset B	1.00	2	0.5

Define the price vector $\mathbf{P} = \begin{bmatrix} 1.65 \\ 1.00 \end{bmatrix}$ and the cashflow matrix $\mathbf{X} = \begin{bmatrix} 3 & 1 \\ 2 & 0.5 \end{bmatrix}$

The goal is to create state price assets, for use as building blocks to price any other asset. The required payoffs from the state price assets are

Required state price payoffs	Current Price	Boom state payoff	Bust state payoff
Boom state asset	?	1	0
Bust state asset	?	0	1

The state price vector $\mathbf{\Psi} = \begin{bmatrix} ? \\ ? \end{bmatrix}$ is unknown.

A replicating portfolio from the two traded assets can be constructed for both state assets. The current price of the state asset will be the current price of the replicating portfolio of traded assets.

$$\mathbf{X}^{-1} = \begin{bmatrix} -1 & 2 \\ 4 & -6 \end{bmatrix} \text{ so } \mathbf{X}^{-1}\mathbf{P} = \begin{bmatrix} -1 & 2 \\ 4 & -6 \end{bmatrix} \begin{bmatrix} 1.65 \\ 1.00 \end{bmatrix} = \begin{bmatrix} 0.35 \\ 0.60 \end{bmatrix} = \mathbf{\Psi}$$

Portfolio holdings	Current Price	Units held of asset A	Units held of asset B
Boom state asset	0.35	-1	2
Bust state asset	0.60	4	-6

The state prices are now defined and can be used to check the price of asset A.

Replication	State Price	Asset A pays	Building block calculation
Boom state	0.35	3	$0.35 \times 3 = 1.05$
Bust state	0.60	1	$0.60 \times 1 = 0.60$

The price of asset A is the sum of the value of the two states $= 1.05 + 0.60 = 1.65$. Note how easy it is now to price an asset. All that is needed is to multiply the payoff in each state by the corresponding state price asset, then total all these to get the current asset price.

In this example there was no risk free bond, but because the state prices are known, one can be constructed. The current price of a bond paying 1 in all states is $0.35 + 0.60 = 0.95$, constructed by holding 3 units of asset A and selling 4 units of asset B [3]. This building block approach can be used to price any asset, where the only requirement is that cashflows of the asset are known in all states. At no point does the probability of either state occurring enter into the calculation. Probabilities are not required because the market is complete, with as many assets as states.

3.1.6 Incomplete markets

If there are as many (linearly independent) assets as states then the market is complete. In a complete market all contingent claims can be bought. These contingent claims do not necessarily have to be explicitly traded, but there needs to be enough securities to span and synthesize all contingent claims.

If there are fewer states than assets then redundant assets exist which, if incorrectly priced, allow for arbitrage opportunities. To take advantage of these opportunities, pick as many assets as there are states, calculate the state prices, then use these to identify arbitrage opportunities in the redundant assets.

If there are more states than assets then the market is incomplete and a different approach is required (this also happens when state space is continuous). In that situation there are an infinite number of possible combinations of state prices that match the cashflows and prices.

Numerical example Adding an extra state to the previous complete market example (making the market incomplete) gives

Price/Payoff matrix	Market Price	Boom state payoff	Bust state payoff	Bad state payoff
Asset A	1.65	3	1	0.5
Asset B	1.00	2	0.5	0.75

[3] This was why a discount rate of 0.95 was chosen in the in appendix A.1

To price Assets A and B correctly it is required that

$$3\Psi_{Boom} + \Psi_{Bust} + 0.5\Psi_{Bad} = 1.65$$
$$2\Psi_{Boom} + 0.5\Psi_{Bust} + 0.75\Psi_{Bad} = 1.00$$

These can be rearranged to define the boom and bust state prices in terms of the bad state price.

$$\Psi_{Boom} = 0.35 - \Psi_{Bad}$$
$$\Psi_{Bust} = 0.60 + 2.5\Psi_{Bad}$$

To avoid arbitrage all state prices must be positive. So in this example the bad state price must be between $0 < \Psi_{Bad} < 0.35$. If positive interest rates are required then an extra bound of $\Psi_{Bad} < 0.02$ is needed, ensuring the sum of the state prices is less than 1.

Although some bounds on the state prices have been established, there are still an infinite number of possible prices for the state assets. State assets cannot give a unique price for other (non-traded) assets so probabilities need to be introduced to progress further.

3.2 Stochastic Discount Factors

In an incomplete market probabilities have to be introduced to complete the market and state prices can become unwieldy in multi-period, multi-asset settings. To simplify the mathematics, define the stochastic discount factor as

$$m(\omega) = \frac{State\ Price}{Probability} = \frac{\Psi(\omega)}{\Pr(\omega)}. \tag{3}$$

SDFs are only well defined for states that have a positive probability of occurring. They can be inserted into the state price equation (eq: 1) after multiplying and dividing by probabilities

$$price(x) = \sum_{\omega} \Psi(\omega)\, x(\omega) = \sum_{\omega} \Pr(\omega)\frac{\Psi(\omega)}{\Pr(\omega)}x(\omega) = \sum_{\omega} \Pr(\omega)\, m(\omega)\, x(\omega) = E[mx] \tag{4}$$

where $\Pr(\omega)$ is the probability of state ω occurring. This is the basic pricing equation. which is mainly just a change in notation of eq: 1. The definition of the SDF looks artificial in a discrete state complete market environment, but expected values are easier to calculate than summations and generalize better to continuous time and space. The underlying state assets are taken care of by design.

Typically there are many possible positive SDFs that can be modelled (as markets are rarely complete), which implies that there are many contingent claim economies consistent with observations and preferences. Introducing probabilities fixes values for the state prices, completing the market.

From this point onwards this paper focuses on SDFs. State assets are largely ignored, but it is useful to remember that the theory underlying the SDF is based on state prices.

3.2.1 Introducing time

Now the time dimension is added to the notation. Every sample ω that is realised produces a different SDF process $m_t(\omega)$, the value of which varies over time t. Similarly the cashflows $x_t(\omega)$ can depend on t as well as on state ω. The pricing equation is extended to assets with cashflows in multiple periods by summing over time as well as states. At time 0 it can be written that

$$price(x) = x_0 = \sum_{\omega}\sum_{t} \Pr_t(\omega)\, m_t(\omega)\, x_t(\omega) = \sum_{t} E[m_t x_t]$$

If the current price of an asset is considered to be just another cashflow then the notation can be simplified. By dropping the ω notation and filtering up to time s, the value of a future cashflow at time t the conditional basic pricing equation can be written as

$$x_s = \frac{E_s\left[m_t x_t\right]}{m_s} \tag{5}$$

where $E_s\left[\,\right]$ is the expected value, given all information up until time s.

Notice that rearranging the equation gives $m_s x_s = E_s\left[m_t x_t\right]$ which is a martingale. An SDF has the property that the deflated price of any cashflow is a martingale. This property is important as it avoids arbitrage in continuous time (see section 3.3.2).

Multiple cashflows may also be priced, so if m is an SDF then the price at time s is $x_s = \frac{1}{m_s}E_s\left[\sum_{j=s+1}^{t} m_j x_j\right]$ is true for all s.

The end time t can be increased by 1 to price a $t+1$ set of cashflows. By induction this can extend the expectation to be over an infinite sum $x_s = \frac{1}{m_s}E\left[\sum_{j=s+1}^{\infty} m_j x_j\right]$, but restrictions on m are required to ensure that the summation is well defined. Going from a t period to ∞ period requires that $\lim_{h\to\infty} E_t\left[m_{t+h}x_{t+h}\right] \to 0$ sufficiently quickly, ruling out bubbles in prices. For the price to be finite, the received cashflows must also be finite. To ensure that cashflows are finite the condition that the price process x_t is mean-summable can be imposed. A process x is mean-summable if $E\left(\sum_{t=0}^{\infty}|x_t|\right) < \infty$. If m is also mean-summable then the dominated convergence theorem implies that $E\left(\sum_{t=0}^{\infty} m_t x_t\right)$ is well defined and finite, therefore the price is finite. To avoid arbitrage all state prices must be positive (see section 3.1.4). Probabilities are non-negative by definition, implying that SDFs must be strictly positive to avoid arbitrage. Price is only dependent on the ratio of the SDF at two time points, so m_0 can be set at any initial value. The standard convention is to define the initial value of the SDF as $m_0 = 1$.

3.2.2 Constructing the SDF in a finite state space

The following section describes how to calculate SDFs whether the market is complete or not. The method is slightly more complicated than the state price calculation, but it does show the link to the basic pricing equation (eq: 4).

Define cashflow space $\underline{C}$ as the set of all cashflows investors can purchase, or the subset of tradeable cashflows used in a particular study. The cashflow space consists of all portfolios or combinations of original cashflows $\left\{\theta^T\mathbf{X}\right\} \in \underline{C}$ where θ is a vector of portfolio weights and $\mathbf{X}$ is the cashflow matrix.

Complete markets Define a specific example of the cashflow space $\underline{C}$ to be generated by S cashflows (S can be thought of as the number of states). This allows the cashflows from an asset i to be organized as a vector $\mathbf{X}_i = \left[x_{i1}, x_{i2}, ..., x_{iS}\right]^T \in \underline{C}$. Similarly, in a complete market there will be $n = S$ assets, so the current prices of the assets can be arranged as a vector $\mathbf{P} = \left[price\left(x_1\right), price\left(x_2\right), ..., price\left(x_n\right)\right]^T$ that spans the cashflow space. Define an SDF vector $\mathbf{M}$ to be $\mathbf{M} = \theta^T\mathbf{X}$ where θ is a vector of portfolio weightings. As the SDF is essentially a portfolio of assets, it is also in the cashflow space $\underline{C}$. The weighting θ is to be constructed so that $\mathbf{P} = \mathbf{X}\mathbf{M}^{\mathbf{T}}$ allowing $\mathbf{M}$ to price all assets. In a complete market the weighting can be constructed without using probabilities. For $\mathbf{P} = \mathbf{X}\mathbf{M}^{\mathbf{T}} = \mathbf{X}\mathbf{X}^T\theta$ to be true the SDF weighting θ needs to be set as

$$\theta = \left(\mathbf{X}\mathbf{X}^T\right)^{-1}\mathbf{P} \tag{6}$$

which, in turn, defines the SDF.

If the market is complete then, for each state, the value of the SDF is the same as the price of the state asset. Only in the complete market case is the SDF unique.

Incomplete markets Incomplete markets exist when contingent claim assets cannot be bought for all states, so $n < S$. Expectations have to be introduced to define the state prices and complete the market. The pricing equation becomes $\mathbf{P} = E\left[\mathbf{X}\mathbf{M}^{\mathbf{T}}\right] = E\left[\mathbf{X}\mathbf{X}^{T}\theta\right]$. Reversing the pricing equation gives a weighting of $\theta = E\left[\mathbf{X}\mathbf{X}^{T}\right]^{-1}\mathbf{P}$, which is used to define the SDF $\mathbf{M} = \theta^{T}\mathbf{X}$. If $E\left[\mathbf{X}\mathbf{X}^{T}\right]^{-1}$ exists and is non-singular then θ exists and is unique, leaving

$$\mathbf{M} = \mathbf{P}^{T}E\left[\mathbf{X}\mathbf{X}^{T}\right]^{-1}\mathbf{X} \tag{7}$$

as the SDF. This process finds an SDF vector that prices all assets, but it does not guarantee that all values in the vector are strictly positive, a requirement for arbitrage opportunities not to exist.

In an incomplete market there are an infinite number of possible SDFs, corresponding to the infinite choice of probability weightings. Picking an SDF by choosing a set of probability weightings defines the pricing of all state assets, completing the market.

Numerical example Taking the previous incomplete market example and assigning probabilities to each state gives

Price/Payoff matrix	Market Price	Boom state payoff	Bust state payoff	Bad state payoff
Real World Probability		10%	50%	40%
Asset A	1.65	3	1	0.5
Asset B	1.00	2	0.5	0.75

Take expectations of the cashflows $\mathbf{X}\,\Pr\mathbf{I} = E\left[\mathbf{X}\right] = \begin{bmatrix} 0.3 & 0.5 & 0.2 \\ 0.2 & 0.25 & 0.3 \end{bmatrix}$ so $E\left[\mathbf{X}\right]\mathbf{X}^{T} = E\left[\mathbf{X}\mathbf{X}^{T}\right] = \begin{bmatrix} 1.5 & 1 \\ 1 & 0.75 \end{bmatrix}$, inverting gives $E\left[\mathbf{X}\mathbf{X}^{T}\right]^{-1} = \begin{bmatrix} 6 & -8 \\ -8 & 12 \end{bmatrix}$. Multiplying by prices $\mathbf{P} = \begin{bmatrix} 1.65 \\ 1.00 \end{bmatrix}$ and cashflows leaves the SDF vector $\mathbf{M} = \mathbf{P}^{T}E\left[\mathbf{X}\mathbf{X}^{T}\right]^{-1}\mathbf{X} = \begin{bmatrix} 3.3 & 1.3 & 0.05 \end{bmatrix}$ and a corresponding state price vector of $\mathbf{\Psi} = \begin{bmatrix} 0.33 & 0.65 & 0.02 \end{bmatrix}$

Different probabilities will produce different state prices. The more likely a state is, the higher the price of the corresponding state asset. Different state prices leads to different SDFs, giving different answers for non-traded cashflows.

3.3 Absence of arbitrage

Arbitrage opportunities offer a positive probability of profit for no cost. There are various ways to ensure absence of arbitrage, depending on the model being used. The following definition may be used when SDFs are applied in finite, discrete state space and time.

Define a set of cashflows x and a pricing function $price\left(x\right)$ to be in cashflow space $\underline{C}$. If the cashflow $x \geq 0$ in all states and in at least one state $x > 0$ then there are no arbitrage opportunities if $price\left(x\right) > 0$. No arbitrage implies that prices of any positive payoffs are also positive. It implies that you cannot create a portfolio for free that might payoff positively but will certainty never make a loss.

Absence of arbitrage also implies that if one asset's cashflows are always greater than those of another asset then its price must be higher. So if $x \geq y$ for all states, and there is at least one state where $x > y$

with positive probability, then $price\left(x\right) > price\left(y\right)$. If $price\left(x\right) < price\left(y\right)$ were true then you would sell short asset y to buy asset x, guaranteeing a positive return.

Absence of arbitrage is equivalent to existence of a positive SDF $m > 0$, and having $m > 0$ implies no arbitrage. But it is not necessarily true that all m satisfying $price\left(x\right) = E\left[mx\right]$ are positive. The existence of a positive SDF does not mean that $price\left(x\right)$ is defined for payoffs $\notin \underline{C}$. For example, if an SDF is defined for a set of assets then it is unlikely that it can be extended to a larger set of assets without implying arbitrage opportunities.

A complete market with no arbitrage implies that there exists a unique $m > 0$ such that $price\left(x\right) = E\left[mx\right]$, implying no arbitrage. In an incomplete market, choosing a unique $m > 0$ completes the market and no arbitrage is assured. If a model is built using positive SDFs then it is guaranteed to be arbitrage free.

3.3.1 SDF assumptions

Assumption 1 - Free portfolio formation

$$x_1, x_2 \in \underline{C} \Rightarrow ax_1 + bx_2 \in \underline{C} \tag{8}$$

for any real a, b where $\underline{C}$ is the space of cashflows.

A portfolio can be constructed from any weighting of assets. This assumption allows investors to trade the underlying assets so that payoffs from any asset can be matched by a portfolio of contingent claim assets. If there is free portfolio formation then it is possible for an SDF to exist. If investors can form a portfolio then it will have a price.

Assumption 2 - Law of one price

$$price\left(ax_1 + bx_2\right) = a \cdot price\left(x_1\right) + b \cdot price\left(x_2\right) \tag{9}$$

The law of one price means that the portfolio has the same value as its contents, so if $x = y + z$ then $E\left[mx\right] = E\left[my\right] + E\left[mz\right]$. This allows all cashflows to be priced by $price\left(x\right) = E\left[mx\right]$. Supply and demand will force the price of the portfolio to be equal to the price of its constituent assets, leading to the assumption of a common price. Existence of an SDF implies the law of one price, but the reverse is also true. The law of one price implies existence of (at least one) SDF.

Combining the above Given free portfolio formation (eq: 8) and the law of one price (eq: 9) then there exists an SDF $m > 0 \in \underline{C}$ such that $price\left(x\right) = E\left(mx\right)$ is unique for all $x \in \underline{C}$ if and only if there are no arbitrage opportunities. Together, both assumptions imply the payoff 0 exists and has a price of 0. SDFs are based solely on the above two assumptions, avoiding other assumptions such as utility functions and complete markets.

3.3.2 Continuous state space

In a complete market, every possible financial claim has a price. In most cases, when the state space is continuous the market is incomplete. This is because there are more possible states (infinite) than there are assets (finite). One role of modelling is to supply market consistent prices where no market price exists. This creates a complete market from an incomplete one and is referred to as completing the market. To complete a market in a continuous state space, probabilities are assigned to each state, and the easiest way to assign probabilities to states is to assume a probability distribution[4]. Completing a market is also equivalent

[4]For example, in the Black Scholes model the market is completed by assuming that asset returns are log-normally distributed.

to picking one SDF from all possible SDFs. Once an SDF has been chosen it can be used to price any future cash flow.

The continuous time and space derivation of the SDF method is different to the discrete case, but ends up with an analogous equation. [Harrison & Pliska, 1981] showed that absence of arbitrage involving asset y is equivalent to the existence of an intrinsic discount process β that makes the process βy into a martingale. The intrinsic discount process β_t generates prices from payoffs, such that the market value at time s, of stochastic cashflow y, due at time $t > s$, is given by

$$y_s = \frac{1}{\beta_s} E_s \left[\beta_t y_t\right] \tag{10}$$

where E_s denotes expectation based on the filtration up to time s.

The goal in constructing the SDF is to ensure that all deflated processes of assets are martingales. Define the current deflated price of an asset to be $y_s = m_s x_s$ and the future deflated price to be $y_t = m_t x_t$. For there to be no arbitrage, the SDF must satisfy the intrinsic discount process (eq: 10) so

$$m_s x_s = \frac{1}{\beta_s} E_s \left[\beta_t m_t x_t\right] \tag{11}$$

Only the ratio of the SDF at two timepoints affects the price x_s. Multiplying the SDF process by a constant value does not change the pricing model, so there is a degree of freedom in the choice of the initial value of the SDF, m_0. The scaling factor chosen to make models easy to build is a multiple of the intrinsic discount process and the original SDF $\beta_t m_t = \widehat{m}_t$. Therefore the revised market value of x at time s is $x_s = \frac{1}{\widehat{m}_s} E_s \left[\widehat{m}_t x_t\right]$. Each SDF $\widehat{m}$ now acts as a state price. Dropping the hat from the notation gives the basic pricing equation in continuous time

$$m_s x_s = E_s \left[m_t x_t\right] \tag{12}$$

In a continuous state space the state prices and probabilities are not a single value, but a continuum of values, therefore the SDF has to be redefined in terms of densities

$$m(\omega) = \frac{State\ Price\ Density}{Probability\ Density} = \frac{\Psi(\omega)}{\Pr(\omega)}$$

Note that in (possibly multidimensional) continuous space the expectation is calculated by integrating with respect to the state ω over all possible states A where $\omega \in A$ so

$$E[mx] = \int_A \Pr(\omega)\, m(\omega)\, x(\omega)\, d\omega$$

3.4 Risk, return and correlation

3.4.1 Risk free assets

An asset is risk free if the payoff is the same across all states. Define $R^f_{s:t}$ as the risk free rate between s and t, so paying $\frac{1}{R^f_{s:t}}$ at time s ensures that an amount of 1 will be returned at time t.

Using the conditional basic pricing equation (eq: 5), if x is the risk free asset then if $x_s = 1$ and $x_t = R^f_{s:t}$ then $m_s 1 = E_s \left[m_t R^f_{s:t}\right] = E_s\left[m_t\right] R^f_{s:t}$ giving

$$\frac{1}{R^f_{s:t}} = \frac{E_s\left[m_t\right]}{m_s} \tag{13}$$

This equation holds in continuous frameworks as well as in discrete ones. It strongly ties the SDF to the bond model, as the expectation of the SDF process must match the pricing given by the bond model. The SDF framework only places minor restrictions on the choice of bond model, which must be arbitrage free and have prices and payoffs defined in the same currency as the SDF.

3.4.2 Covariance pricing model

Define the return between time s and t by $R_{s:t} = \frac{x_t}{x_s}$ and the SDF at time t to be m_t. Consider expected returns on investment using the basic pricing equation (eq: 4). Using the definition of covariance, the equation can be rearranged to

$$m_s 1 = E_s\left[m_t R_{s:t}\right] = E\left[m_t\right] E\left[R_{s:t}\right] + Cov\left(m_t, R_{s:t}\right) \tag{14}$$

Using the SDF's risk free rate identity (eq: 13) the equation can be rearranged to

$$\frac{E_s\left[R_{s:t}\right]}{R^f_{s:t}} = 1 - Cov\left(R_{s:t}, \frac{m_t}{m_s}\right) \tag{15}$$

The risk premium is related to the covariance of the asset return with the change in SDF value. The equation can also be written in terms of prices

$$Cov\left(m_t, x_t\right) = E\left[m_t x_t\right] - E\left[m_t\right] E\left[x_t\right] = x_0 - \frac{E\left[x_t\right]}{R^f_{0:t}}$$

so

$$x_0 = \frac{E\left[x_t\right]}{R_f} + Cov\left(m, x_t\right) \tag{16}$$

This is the present value equation for a risk free asset, plus a risk adjustment. The risk adjustment is the covariance of the asset with the SDF.

A positive risk premium is obtained for taking risks that are negatively correlated with changes in the SDF. e.g. a positive risk premium means that the asset is expected to fall when the SDF rises (and vice versa).

Note that in eq: 16, zero correlation with the SDF gives $x_s = \frac{E[x_t]}{R_f}$ which is the expected value of the cashflow (in real world terms) discounted at the risk free rate, independent of the risk of x_t. SDFs only assign value to risks correlated with themselves.

Excess returns Defining excess return as returns above the risk free rate $\frac{R^i_{s:t}}{R^f_{s:t}} = {}^e R^i_{s:t}$ and using the basic pricing equation, rearranged in terms of returns (eq: 14) combined with the risk free property (eq: 13) gives,

$$m_s = m_s E_s\left[{}^e R^i_{s:t}\right] + Cov\left(m_t, R^i_{s:t}\right) \tag{17}$$

So defining $R^m_{s:t} = \frac{m_t}{m_s}$ means

$$1 = E_s\left[{}^e R^i_{s:t}\right] + Cov\left(R^m_{s:t}, R^i_{s:t}\right)$$

This explicit relationship, between expected excess return and the correlation between change in the SDF and change in price, is how the present value of future prices is defined. The risk premium of an asset is controlled by its correlation with the SDF.

For example, if equities are assumed to return more than the risk free rate then an appropriate correlation with the SDF must be assumed, to ensure that the present value for the equities equals the expected deflated future value. Conversely, any assets uncorrelated with the SDF will be discounted at the risk free rate. The covariance structure does not have to be constant over time, and can even be conditional on state. Any stochastic price process may be implemented, so long as the discounted asset process is a martingale.

3.4.3 Risk premiums

Risk premiums exist because a risk averse investor is more concerned about a drop in their wealth than an increase in wealth of the same amount. If the investor is fully diversified then they are exposed to all risks in the economy, but only by a tiny amount to the risks in each asset. Risks that are not correlated or common to all assets are of little concern, because not all of these risks will occur at the same time. However, there will be concern over common risks across the portfolio as these risks are correlated with one another and are able to affect the investor's overall wealth. Extra financial compensation is required to bear the risk that the correlated assets may simultaneously drop in value, so a risk premium is demanded.

Rather than consider all correlations between all assets, it is simpler to reference the correlated risks to a particular weighting of all assets called the market portfolio (see section 4.1.1).

3.4.4 Pricing multiple assets

If contingent claims to many states of nature are not available then the SDF m is not unique. Unless the market is complete, there are an *infinite* number of random variables satisfying the basic pricing equation (eq: 4). When calibrating an SDF model there will typically be more than one SDF to choose from that will price all traded investments at market prices.

If there exists an SDF m such that $price\,(x) = E\,[mx]$ then the equation $price\,(x) = E\,[(m+\varepsilon)\,x]$ generates all additional discount factors, for any ε orthogonal to x (i.e. $E\,[x\varepsilon] = 0$). So the SDF m can be represented as $m = m^* + \varepsilon$ where $E\,[\varepsilon x] = 0$.

The breakdown of an SDF into a correlated component and an uncorrelated component allows a single SDF to price multiple assets. Whilst the correlation of the asset with the m^* part of the SDF prices one asset, the ε part of the SDF prices other assets. Different assets will have different breakdowns between the correlated and uncorrelated parts of the SDF.

3.5 Model building

There are many approaches to defining the SDF process, but the approach introduced in this section is simpler than those in the literature, whilst still being generic. The requirements placed on a price process for use within the SDF model are discussed. Finally a log-normal, mean reverting example is given, employing the new, simple SDF modelling approach.

3.5.1 General model

It is important to be able to control the expected future value of the SDF as the mean must be the same as the price $\frac{1}{R^f_{s:t}}$ of a risk free zero coupon bond with a payoff of 1, as given in eq: 13. This constraint can be rewritten as $E_s\,[m_t] = \frac{m_s}{R^f_{s:t}}$. An easy way to satisfy the constraint is to break the SDF into two independent parts, building part of the SDF out of a bond model.

$$m_t = m_s\,(Bond\ Model)_{s:t} \times (Random\ Innovation)_{s:t} \tag{18}$$

where $E\,[(Random\ Innovation)_{s:t}] = 1$ so that

$$E_s\,[m_t] = m_s\,(Bond\ Model)_{s:t} = \frac{m_s}{R^f_{s:t}}$$

The SDF must not become negative. The bond model will not produce negative prices (if it is arbitrage free) so long as the starting value for the SDF is greater than zero $m_0 > 0$ (convention sets $m_0 = 1$) and the random innovation is strictly positive. The random innovation element may be modelled without reference to the bond model, although correlation is required if the bond model is not risk free. An easy way to ensure positivity is to model exponential processes.

$$(Random\ Innovation)_{s:t} = \exp\left[(Random\ Process)_{s:t}\right]$$

Any (arbitrage free) bond model[5] may be used, although the bond model must be able to give a price for the whole term structure. For example, if monthly time steps are generated via Monte Carlo simulation then the bond model must be able to price a bond maturing in one month's time. A stochastic interest rate will add noise to the SDF (in comparison to deterministic interest rates) increasing the number of Monte Carlo runs required for prices to converge.

If a new asset is added to an existing SDF model then the SDF process does not need to be modified, so long as the correlation between the SDF and the new asset is correct. Correlations can be created in a number of other ways. Factor modelling can be used to create new random processes based on the value of assets already projected in the SDF model. Alternatively, a new random variable common to the SDF and the asset can be added to the SDF. This approach is not ideal as it increases the volatility of the SDF.

3.5.2 Simulation of the SDF

Modelling the SDF by using a bond model multiplied by a random innovation is relatively easy to do. To introduce SDFs into an existing Monte Carlo model, the only addition required is a new stochastic process. To generate the future SDF value:-

- take the current SDF value,
- multiply by the bond price of a zero coupon bond, paying 1 at the end of the next time step,
- multiply by a suitably correlated, strictly positive, random innovation that has a mean value of 1.

At this point, few restrictions have been placed on the price process (it is implicitly assumed that the expectations of all processes are finite). All distributions, seasonal factors, mean reversions etc. are allowed within this framework. The only requirement is that it must be possible to control the covariances between the price processes and the SDF.

In a Monte Carlo implementation, the covariance structure only needs to be calculated as far as the next time step. This simplifies the projection of more complicated price processes (like time changing parameters).

There is only one SDF process per model. This means that the pricing of derivatives should be calibrated through the covariance between the SDF and the asset, not through the choice of SDF distribution. Correlating processes to the SDF can be made numerically simple and flexible by using an analytical equation to control the risk premium through the correlation matrix. The copula approach allows many different types of stochastic processes to be combined but analytical tractability is sacrificed.

3.5.3 Multivariate log-normal model

The following multivariate log-normal model is original to this paper. It can be simulated easily, allows a forward curve, and its outcomes can be checked analytically. Asset risk premiums and covariances of the SDF to all prices are defined through the multivariate normal distribution. Importantly, an explicit link

[5]The literature on bond modelling is extensive and well researched. This paper does not extend the bond model research area.

between the risk premium and the correlation parameter in the multivariate normal covariance matrix is given. The model is extended to mean reverting assets in section 3.5.4.

The log-normal spot price $S(t)$ at time t for asset i can be written as

$$S(t)\,|S(s) = F(s,t)\exp\left[-\frac{1}{2}\sigma_i^2(t-s) + \varepsilon_i\sigma_i\sqrt{t-s}\right] \tag{19}$$

where $\varepsilon_i \sim MultivariateNormal(0,1)$ and $F(s,t)$ is the forward price for asset i between time s and t. Now we guess that an SDF exists of the form

$$m_t = m_s\exp\left[\left(A_{s:t} - \frac{1}{2}\sigma_m^2\right)(t-s) + \varepsilon_m\sigma_m\sqrt{t-s}\right] \tag{20}$$

where $A_{s:t}$ is a function to be defined that depends upon timepoints s and t.

Taking the expectation of the SDF gives $E_s[m_t] = m_s\exp[A_{s:t}(t-s)]$ and using the risk free definition (eq: 13) it can be rearranged, resulting in $A_{s:t} = \frac{-\ln\left(R^f_{s:t}\right)}{(t-s)}$. So the SDF is defined as

$$m_t = \frac{m_s}{R^f_{s:t}}\exp\left[-\frac{1}{2}\sigma_m^2(t-s) + \varepsilon_m\sigma_m\sqrt{t-s}\right] \tag{21}$$

This completes the process for the SDF leaving only the volatility of the SDF σ_m and its correlations ρ_{im} to define.

Both the price process and the SDF process are log-normal stochastic processes, so there is a closed form representation of covariance

$$Cov(m_t, S(t)) = \frac{m_s}{R^f_{s:t}}(s,t)\,F(s,t)\left[\exp\left(\rho_{im}\sigma_i\sigma_m(t-s)\right) - 1\right] \tag{22}$$

Inserting eq: 22 into the covariance pricing model (eq: 16) and simplifying gives the definition for excess return in terms of correlation with the discount factor as

$$E_s\left[{}^eR^i_{s:t}\right] = \exp\left(-\rho_{im}\sigma_i\sigma_m(t-s)\right) \tag{23}$$

where return is defined as $R_{s:t} = \frac{F(s,t)}{S(s)}$ so excess return is ${}^eR^j_{s:t} = \frac{F(s,t)}{S(s)}\frac{1}{R^f_{s:t}}$

To ensure that $|\rho_{im}| \leq 1$, a minimum variance for the SDF is defined as

$$\sigma_m \geq \max_i\left|\frac{-\ln E_s\left[{}^eR^i_{s:t}\right]}{\sigma_i(t-s)}\right| \tag{24}$$

This means all assets priced by the SDF are within the efficient frontier wedge (see Fig 4.1.1).

Finally the correlations between each asset and the SDF are set as

$$\rho_{im} = \frac{-\ln E_s\left[{}^eR^i_{s:t}\right]}{\sigma_i\sigma_m(t-s)} = \frac{-\ln E_s\left[{}^eR^i_{annual}\right]}{\sigma_i\sigma_m}. \tag{25}$$

This is an instantaneous risk premium, defined in annual terms following common convention, although the period chosen for the timepoints s and t is arbitrary. Another interpretation is that if the risk premium is time homogeneous then so are the correlations. Note that excess return is solely defined in terms of covariance with the SDF. The forward price of asset i has been canceled out of the equation.

If a constant time step and time homogeneous risk premium is used then only one covariance matrix is required. If a stochastic interest rate is modelled then, for the covariance matrix to remain constant, total return must be modelled as the excess return of an asset on top of the stochastic interest rate.

If the SDF is perfectly negatively correlated with one of the variables, k, such that $\rho_{km} = -1$ then, as the SDF is the reciprocal market portfolio, asset k is the market portfolio (as defined by the model). If all of the variables have $\rho_{im} > -1$ then no asset in the model is the market portfolio, although it exists by proxy within the model as the reciprocal of the SDF.

It is easy to insert this simple form of the SDF framework into existing variance-covariance models as all that is required is the addition of a new log-normal process, which is the same as adding a new asset process. Care must be taken when using the correlation method as it is very easy for a matrix to not be positive definite. Assigning a larger variance to the SDF gives smaller correlations, making a positive definite matrix easier to obtain.

3.5.4 Multivariate mean reverting log-normal model

If a mean reverting term is added to a log-normal price process then the distribution is still log-normal, but the volatility is dependent on the stepsize and the mean reversion parameter, and the expected return is dependent on the spot. [Clewlow & Strickland, 1999] derive the spot price of a mean reverting processes of asset j with mean reverting parameter α as

$$S(t)\,|S(s) = \hat{F}(s,t)\exp\left[-\frac{1}{2}\hat{\sigma}^2_{s:t} + \varepsilon_j\hat{\sigma}_{s:t}\right] \tag{26}$$

where the modified variance parameter is

$$\hat{\sigma}^2_{s:t} = \frac{\sigma_j^2}{2\alpha}\left[1-\exp\left(-2\alpha\,(t-s)\right)\right] \tag{27}$$

and the modified forward rate is

$$\hat{F}(s,t) = F(0,t)\left(\frac{S(s)}{F(0,s)}\right)^{\exp(-\alpha(t-s))}\exp\left(-\frac{\sigma^2}{4\alpha}\left(e^{-\alpha t}\left(e^{2\alpha s}-1\right)\left(e^{-\alpha t}-e^{-\alpha s}\right)\right)\right) \tag{28}$$

This means that

$$Cov\,(m_t, S(t)) = \frac{m_s}{R^f_{s:t}}\hat{F}(s,t)\left[\exp\left(\rho_{jm}\sigma_m\hat{\sigma}_{s:t}\sqrt{(t-s)}\right)-1\right]$$

so

$$E_s\left[{}^eR^j_{s:t}\right] = \exp\left(-\rho_{jm}\sigma_m\hat{\sigma}_{s:t}\sqrt{(t-s)}\right) \tag{29}$$

or

$$\rho_{jm} = \frac{-\ln E_s\left[{}^eR^j_{s:t}\right]}{\sigma_m\hat{\sigma}_{s:t}\sqrt{(t-s)}} = \frac{-\ln E_s\left[{}^eR^j_{annual}\right]}{\sigma_j\sigma_m\sqrt{\frac{[1-\exp(-2\alpha(t-s))]}{2\alpha(t-s)}}} \tag{30}$$

so because $\rho \to 0$ as $t-s \to \infty$ this means that the risk premium of a mean reverting process disappears over time (see appendix A.4). If $t \to s$ or $\alpha \to 0$ then the non-mean reverting version (eq: 25) is recovered, implying that if small enough timesteps are simulated then the correlation matrix will not need to be modified. Over a small timestep the noise term dominates the mean reverting drift of the process. A two asset numerical example of eq: 30 is implemented in section 6 using Monte Carlo.

4 Real Option Valuation

This section introduces a new approach to real options analysis. Using the SDF can simplify calculations and assumptions, whilst allowing models to be generic and flexible.

The real option approaches listed in the literature review take an average value of future cashflows and then discount them to the present. The real option approach used in this paper reorders this process. Future cashflows are **first** discounted to the present and **then** averaged.

subsection 4.1 discusses various approaches, comparing them to the SDF approach in particular. Only the argument against utility functions in section 4.1.3 is original in this valuation section. The second subsection 4.2 is non-mathematical but is important for justifying the practical implementation of the Monte Carlo SDF approach to real options.

4.1 Valuation approaches

4.1.1 Risk adjusted discount rates

A traditional way to calculate the present value of risky cashflows, and therefore price risky assets, is to discount using a risk adjusted rate. The present value of a project is the expected value of future cashflows, discounted at a risk adjusted rate that is equal to the expected return in eq: 35, so

$$price\left(x^{i}\right)=\frac{E\left[x^{i}\right]}{R^{i}}$$

The capital asset pricing model (CAPM) equation can be derived from SDFs. The starting point is the basic pricing equation (eq: 4) expressed in terms of returns

$$1=E\left[mR^{i}\right]=E\left[m\right]E\left[R^{i}\right]+\sigma\left(m\right)\sigma\left(R^{i}\right)\rho_{m,R^{i}}$$

Where $\rho_{m,R^{i}}$ is the correlation between the SDF m and the return on asset i and $\sigma\left(\cdot\right)$ is the standard deviation operator. Using the risk free rate identity (eq: 13) this can be rearranged to

$$E\left[R^{i}\right]-R_{f}=-R_{f}\sigma\left(m\right)\sigma\left(R^{i}\right)\rho_{m,R^{i}} \tag{31}$$

The equation implies an asset with a positive expected risk premium has a negative correlation with the SDF. Correlations are always bounded by $|\rho|\leq 1$ implying a bound on the risk premium

$$\left|E\left[R^{i}\right]-R_{f}\right|\leq R_{f}\sigma\left(m\right)\sigma\left(R^{i}\right)$$

which is represented graphically in fig 4.1.1 (figure is adapted from [Cochrane, 2004]).

Define the market asset $R^{i}\rightarrow R^{p}$ by picking a point on the upper bound line, so that $\rho_{R^{p},m}=-1$. Inserting the market asset into eq: 31 allows the frontier line to be written as

$$R_{f}\sigma\left(m\right)=\frac{E\left[R^{p}\right]-R_{f}}{\sigma\left(R^{p}\right)} \tag{32}$$

relating the Sharpe ratio of the market (also referred to as the slope of the frontier, or the market price of risk) to the SDF volatility. Any point along this efficient frontier can be replicated by a combination of the market asset and a risk free asset. Assets on the line are perfectly correlated, with each other and the SDF.

The slope of the efficient frontier is governed by the volatility of the SDF, placing a constraint on the SDF. It needs to be volatile enough to ensure that all assets available are contained within the wedge. All assets are contained within the wedge because, by definition, there are no assets that are more efficient.

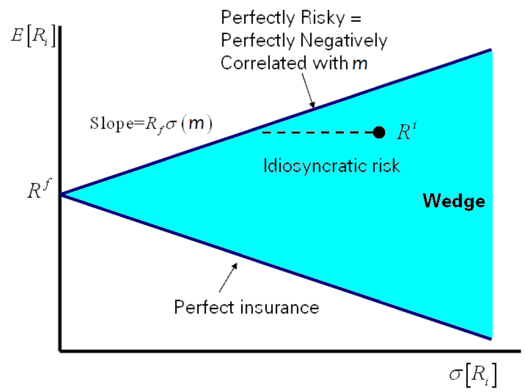

Figure 1: Volatility Wedge

Assets inside the frontier assets are not inferior. Rational investors are happy to hold all assets whether they are on the efficient frontier or not. They are not worried about the idiosyncratic risk of each asset because not all of the idiosyncratic risks will occur at the same time.

An asset with the maximum Sharpe ratio will be perfectly negatively correlated with the SDF. One of these assets has already been defined to be the market asset R^p. As $\rho_{R^p,m} = -1$ there is a direct linear relationship $m = a - bR^p$ for constant a and positive scalar b. This linear relationship shows that $\sigma(m) = b\sigma(R^p)$ leads to

$$\sigma\left(R^i\right)\sigma(m)\,\rho_{R^i,m} = Cov\left(R^i, m\right) = Cov\left(R^i, a - bR^p\right) = -Cov\left(R^i, R^p\right)b \tag{33}$$

After rearranging this becomes

$$\sigma\left(R^i\right)\rho_{R^i,m} = -\frac{Cov\left(R^i, R^p\right)}{\sigma(R^p)} \tag{34}$$

So putting eq: 32 and eq: 34 into eq: 31 gives

$$E\left[R^i\right] = R^f + \beta^i\left(E\left[R^p\right] - R^f\right) \tag{35}$$

where $\beta^i = \frac{Cov\left(R^p, R^i\right)}{V[R^p]}$. This is the capital asset pricing model (CAPM), giving an important conclusion. The SDF can be thought of as the inverse of the market portfolio.

The SDF is perfectly negatively correlated with the market, so its value is high when the market is low and low when the market is high. Essentially, the SDF rewards cashflows that are produced in a bear market more than ones produced in a bull market. Pricing with reference to CAPM and the efficient frontier is often referred to as equilibrium pricing.

The SDF approach does not need the market portfolio to be explicitly defined within the model. This is because the market portfolio and the SDF are equivalent as they are perfectly negatively correlated. Existence of an SDF implies an efficient frontier and an efficient frontier implies that an SDF exists.

4.1.2 Risk neutral valuation

A popular real option valuation approach is to apply risk neutral probabilities (see appendix A.1.2) to price real options. The assumption of risk neutrality, and the use of the risk free discount rate, is justified by identifying a portfolio of traded investments exactly replicating the cashflows of the real option, removing all risks from the asset (also assuming perfect and frictionless markets) making it equally valuable to all investors. Using no arbitrage arguments (see appendix A.2), the value of the option must equal the value of the replicating portfolio.

This approach to real options works well where underlying risks are freely traded (e.g. with commodity based real options), but the risk neutral approach is harder to justify when the risks cannot be hedged in the market. Many assets are not freely traded and there are few very long term traded options available with which to construct hedges.

Risk neutral probabilities Risk neutral probabilities are an artificial construct that, by changing the probabilities of the states occurring, simplify pricing models (see the numerical example in appendix A.1.2). The following derivation of risk neutral probabilities using state prices is based on [Etheridge, 2002], who makes the mathematics easier to follow by using discrete time and space.

Consider a complete market with n securities (as in section 3.1.4). If there is a state price vector Ψ that has been derived from the cashflow matrix $\mathbf{X}$ with n states and price vector $\mathbf{P}$ then a new value can be defined as $\Psi_0 = \sum_{\omega=1}^{n} \Psi(\omega)$. For any state ω, set the risk neutral probability of that state to be $q(\omega) = \frac{\Psi(\omega)}{\Psi_0}$. As $\sum_{\omega=1}^{n} q(\omega) = 1$ and $q(\omega) \geq 0$ the vector $[q(1), \ldots, q(n)]$ is a well defined probability distribution, referred to as the risk neutral distribution. Using eq: 1 the price of an arbitrary security with a cashflow x can be written as

$$\frac{price\,(x)}{\Psi_0} = \sum_{\omega=1}^{n} q(\omega)\, x(\omega) = E^Q[x] \tag{36}$$

These q probabilities satisfy the requirements to be a state price, as $1 \geq q(\omega) > 0$.

If there exists an asset B with payoffs of 1 in all states, so $B(\omega) = 1$ for all ω, then using 36 gives $\frac{price(B)}{\Psi_0} = \sum_{\omega=1}^{n} q(\omega) = 1$ which defines $price\,(B) = \Psi_0 = \frac{1}{R^f}$ as the discount rate on riskless borrowing.

The normalized prices can be viewed as the expected payoff under a specially chosen set of risk neutral probabilities.

$$price\,(x) = \sum_{\omega=1}^{n} \Psi(\omega)\, x(\omega) = \frac{1}{R^f} \sum_{\omega=1}^{n} q(\omega)\, x(\omega) = \frac{E^Q[x]}{R^f} \tag{37}$$

where E^Q uses the risk neutral probability measure Q. Any security x can now be written as $price\,(x) = \frac{1}{R^f} E^Q[x]$ showing the security's price to be its discounted expected payoff under artificially constructed probabilities.

Risk averse investors give greater value to cashflows received in bad states than good states, therefore the risk neutral probability measure Q will give greater weight to bad states as their state prices will be larger. Simulating under a risk neutral measure will cause systematically bad events to occur more often than good events.

Risk neutrality and the SDF Risk neutral probabilities are actually a special case of the SDF. This can be seen by comparing eq: 4 to eq: 37, so

$$E^P[mx] = \frac{1}{R_f} E^Q[x] \tag{38}$$

where probability measure P is the real world measure. Inspecting eq: 38 and conditioning on each state it can be seen that the SDF can also be defined as

$$m(\omega) = \frac{q(\omega)}{p(\omega)} \frac{1}{R^f} \tag{39}$$

showing that the SDF is the ratio of the risk neutral measure to the real world measure, discounted at the risk free rate[6].

State prices provide the link between risk neutral and real world probabilities. Risk neutral probabilities can be converted to state prices which can be converted to real world probabilities (and vice versa). Risk neutral pricing can be derived from SDFs, so anyone comfortable holding unhedged derivatives, priced using risk neutral probabilities should be happy using SDFs.

4.1.3 Utility functions

Utility functions are used to define risk preferences of investors, which are then used to calculate the risk premium on a real option. Using utility functions replaces the objective of maximizing shareholder value with an objective of maximizing a risk preference.

The approach is not suitable if shareholders are rational and seek to diversify their assets. If utility functions alone were used to price assets then shareholders would be happy to accept any price and would vary their shareholdings to accommodate the risk.

To demonstrate this, define an investor to have the utility function $U(x)$ where x is the wealth. The expected change in utility can be measured by an approximation using the power series expansion

$$\begin{array}{c} \frac{E[U(x+\varepsilon)]}{U(x)} = \frac{E\left[U(x)+\varepsilon U'(x)+\frac{1}{2}U''(x)\varepsilon^2\right]+o\left(\varepsilon^3\right)}{U(x)} \\ \simeq 1+\frac{U'(x)}{U(x)}E[\varepsilon]+\frac{1}{2}\frac{U''(x)}{U(x)}\left(E[\varepsilon]^2+V[\varepsilon]\right) \end{array} \tag{40}$$

where ε represents an uncertain future cashflow.

The investor can invest a proportion y in the risky security and the remaining $1-y$ in the risk free security. So the random component can be broken down as $\varepsilon = \varepsilon^f + y\left(\varepsilon^i - \varepsilon^f\right)$ giving $E[\varepsilon] = E\left[\varepsilon^f\right] + yE\left[\left(\varepsilon^i - \varepsilon^f\right)\right]$ and $V[\varepsilon] = V\left[\varepsilon^f + y\left(\varepsilon^i - \varepsilon^f\right)\right] = y^2 V\left[\varepsilon^i\right]$ because the risk free asset has no variance $V\left[\varepsilon^f\right] = 0$.

To simplify the notation assume $\varepsilon^f = 0$ and scale current utility so that $U(x) = 1$, giving the quadratic $E[U(x+\varepsilon)] = ay^2 + by + 1$ where $a = \frac{1}{2}U''(x)V\left[\varepsilon^i\right]$ and $b = \left(U'(x) + \frac{1}{2}U''(x)E\left[\varepsilon^i\right]\right)E\left[\varepsilon^i\right]$. Solving the quadratic for y and differentiating shows that utility is maximised when

$$y = \frac{E\left[\varepsilon^i\right]}{V\left[\varepsilon^i\right]}\left(-\frac{E\left[\varepsilon^i\right]}{2} - \frac{U'(x)}{U''(x)}\right)$$

The greater the return, the higher the proportion y. The larger the variance of the security, the smaller the proportion. Due to the non-linearity in the pricing the investor will happily accept any variance for the underlying security by varying the proportion invested accordingly.

If the investment cannot be broken down into liquid sub-units, or it has a small number of owners then utility functions are applicable. But if a company has shares traded openly then each shareholder can individually increase or decrease their holdings of the company (and therefore the underlying real option) until their individual risk preference is satisfied. This means that utility functions should not be applied to decision making within publicly traded companies.

If $U''(x) = 0$ then variance is removed from eq: 40 which then becomes linear. This utility function is a special case where the investor is said to be risk neutral.

[6] This relationship is analogous to the Radon-Nikodym derivative.

4.1.4 Real world probabilities

Appendix A.1.1 gives a numerical example showing the potential pitfalls of pricing using real world probabilities.

When constructing SDFs, the state price can be divided by any chosen probability measure, so a natural choice is to use the probabilities that actually occur. The SDF uses the same real world probability measure as the cashflow payoffs, so there is a consistent probability measure across all processes. Using SDFs with real world probabilities allows all assets and derivatives to be priced but also allows price projections to include risk premiums, giving realistic prices which allow other risk measures to be calculated.

4.1.5 Certainty equivalent values

The certainty equivalent approach requires an adjustment CE^i to be made to each (real world) expected cashflow $E\left[x^i\right]$, so that it may be discounted at the risk free rate.

$$price\left(x^i\right) = \frac{E\left[x^i\right] - CE^i}{R^f} \tag{41}$$

matching eq: 16 to eq: 41 implies that

$$CE^i = -R^f Cov\left(m, x^i\right)$$

or rearranging using the frontier equation (eq: 32) and the linear relationship between the SDF and the market (eq: 33) gives

$$CE^i = Cov\left(R^p, x^i\right) \frac{E\left[R^p\right] - R^f}{\sigma\left(R^p\right)^2} \tag{42}$$

which is the more common form of the certainty equivalent adjustment. This link between the SDF and certainty equivalent cashflows shows that the two approaches are mathematically equivalent.

The certainty equivalent approach is good for simple calculations (e.g. a binomial tree with few timesteps and only one source of risk) and produces transparent, analytic solutions using real world probabilities.

The main drawback in working with certainty equivalent values is that correlation with the market portfolio must be evaluated for every risky cashflow. So for a large project the calculations become burdensome unless major simplifying assumptions are made about both the cashflow profile and the underlying risks.

4.2 SDFs and real options

A project (or portfolio of projects) containing real options can be priced using the covariance model in section 3.5. The cashflows can be discounted using the SDF to give a present value for the project. The value of each source of flexibility is the change in project value for that decision. Using SDFs, both the underlying real asset and options on that underlying asset are priced simultaneously.

The Monte Carlo SDF approach has various advantages over traditional real option pricing methods; it does not require a separate calculation for each cashflow (certainty equivalent does); it uses real world probability measures (risk neutral does not); and it uses the same discount rates for all cashflows (risk adjusted rates depend on the asset being valued).

Using the covariance SDFs combined with Monte Carlo techniques can provide models that are simple to implement, yet retain accuracy and realism making real options more likely to be employed in practice. If each option is valued using the SDF then the portfolio will be priced without arbitrage and without having to resort to any hedging arguments.

4.2.1 Mathematical equivalence

SDFs will not produce different values to other techniques. Their main advantages are that they are simple to explain and implement, and allow asset prices to be projected under the real world probability measure. The equivalence of SDFs to various other modelling methods can be used to check simple results and explain more complex ones.

4.2.2 Market and subjective data

The Monte Carlo SDF model appropriately incorporates risks within its framework by employing as much information as available, from the market and from experts[7].

Market data comprises of the past history of prices, and current prices of derivative assets. It may also be possible to use traded derivatives to work out the implied probabilities and parameters, as long as the derivatives are sufficiently liquid. If historic price data is available for an asset then SDF correlations can be calibrated using the historical correlation between the asset and the market. A diversified market portfolio, such as the FTSE All Share or S&P 500 may be an appropriate choice of market proxy for SDF correlations. [Breymann et al., 2005] find that the Morgan Stanley Capital Growth World Index (MSCI) is a good proxy for the theoretical growth optimised portfolio ([Platen, 2006]).

The subjective approach assumes that assets and markets exist, but does not identify them. Experts must be relied upon for subjective estimates of statistics. For the SDF this involves defining either the asset's expected risk premium, or its correlation with the market portfolio (and therefore with the SDF, see section 4.1.1).

4.2.3 Systematic and unsystematic risks

SDFs are the ideal method for tackling systematic and unsystematic risks within the same framework. SDFs value assets from the perspective of a fully diversified shareholder. SDFs place greater value on systematic cashflows that occur in poor times than good times. Unsystematic risks are uncorrelated with the general economy, so the SDF discounts their cashflow at the risk free rate (on average).

4.2.4 Complete markets

Valuing an option using the replicating portfolio approach requires a perfect hedge, removing all risk. The market needs to be complete for an option to be perfectly hedged. SDFs do not require a perfect hedge to price an option, so they do not need complete markets. Unfortunately an incomplete market implies the existence of an infinite number of possible SDFs, giving an infinite number of possible prices for the option. The SDF approach offers two solutions to this infinite price problem.

One approach for completing the market is by picking an SDF from the infinite set of possible SDFs, choosing the SDF's functional form. [Black & Scholes, 1973] took this approach to completing the market by choosing Brownian motion for their option pricing model. Choosing an SDF with a low volatility lets the Monte Carlo converge faster.

The other solution, the good deal bound approach, is used to place limits on the price of the residual risk, constraining the total market price of risk below a reasonable value. Again, a lower volatility SDF produces a tighter price bound.

[7] If competitive interactions are required then these have to be implemented as a separate modelling layer, using other techniques, like game theory

4.2.5 Cashflow independence

A common mistake in applying real option analysis is to focus on the options available in the revenues of a project and ignore options on costs. SDFs value cashflows independently of whether they are positive or negative (costs or revenues). The approach can be used at company, project and even transaction level, complementing, rather than replacing other tools.

The same underlying set of asset projections can be produced for many purposes. Specific business problems can be modelled in parallel, only drawing on data from the Monte Carlo SDF model when prices are required. The SDF method only has to be described to management once, allowing their attention to be focused on the more valuable areas of identifying real options and ensuring they are executed correctly.

The mathematical complexities involved in generating the stochastic evolution of asset prices and a correlated SDF process can be hidden away. Only the generated numbers need to be viewed by practitioners. Growth assumptions are offset by the risk premium, so have no effect on present value calculations when the SDF approach is used (with normal distributions), simplifying decision making and removing a major source of contention.

4.2.6 Disadvantages of the SDF

The Monte Carlo SDF approach takes longer to converge to a solution than other approaches, so the approach is recommended when time is not an issue, and real world probabilities are required. Techniques to speed up the Monte Carlo, such as control variates and antithetic random numbers, are still applicable to SDFs.

The SDF is a measuring tool. It does not solve any decision making problems. Other solution techniques, appropriate to the problem under consideration, should be employed alongside the measurements given by the SDF method. Although the Monte Carlo method allows many theories to be tested, optimal strategies are hard to identify because each strategy being tested usually requires its own Monte Carlo run. The SDF method takes slightly more simulations to converge to a stable value than the risk neutral Monte Carlo method, as there is the added variation of the SDF process to average out, but this is a small price to pay for the benefits gained through real world probability modelling.

If risk premiums are not too large and the correlation matrix is positive definite then the SDF volatility can be kept small, reducing the number of simulations required for convergence. Manipulating the correlation matrix so that the volatility of the SDF is kept low, whilst still remaining positive definite, can be a matter of judgement.

SDFs only calculate the value of an asset. A replicating portfolio is not explicitly identified so the portfolio required to hedge a specific real option is generally not known. Other pieces of information, such as the amount of risk created or sensitivities to changes in underlying parameters[8] are not calculated. The risk when using SDF has to be controlled separately.

The Monte Carlo SDF approach is not the answer to everything. Analytic solutions are best found through other methods. Absence of a replicating portfolio or any measurement of parameter sensitivity makes the approach unsuitable for trading applications. Removal of the replicating portfolio means that risks have to be monitored and controlled by other methods. This is explored in section 5.

5 Controlling risk

Most real options approaches rest on the assumption that, in complete markets, the investor will find a traded security that exactly replicates the risk of the option's payoff at any point in time between acquisition and

[8] Referred to as Greeks in financial option pricing

exercise. Since the current price of this replicating portfolio is known, the current value of the option is known. The investor may also use the replicating portfolio to hedge the option, removing risk.

5.1 Hedging

The necessity to hedge a risk provides the largest stumbling block for most real option approaches. The Monte Carlo SDF approach removes the hedging requirement, allowing pricing without resorting to the assumption of a replicating portfolio. Instead SDFs use the correlation with the market, avoiding all difficulties surrounding the tradeability of an asset.

5.1.1 Hedging individual options

It is usually difficult to discover and create a suitable replicating portfolio as the underlying assets in most real options are not traded.

Specific real options hedging by an individual firm in a highly competitive area will have an immediate impact on the action of its competitors, possibly signaling valuable information to the market. Exercising the real option may alter the momentum of a the perfectly matched hedging security. Real options contain many operational and contractual obligations which tend to be inseparable from other real and intangible assets, meaning they may only be sold as part of a package. A real option can be held collectively by many firms, making a sale by an individual company impossible.

Even if a replicating portfolio can be found, it is often not feasible due to substantial costs and imperfect markets. The underlying assets are usually not liquid enough to trade or sell short and relevant futures contracts may not even exist.

The execution strategy of a real option is unlikely to be changed by the purchase of a hedging financial option. The hedge is there to control cashflow volatility. Holding a hedging portfolio places a limit on just how much downside risk a single real option can impose on a firm. But controlling the risk at the project level, rather than the company level, is likely to be ineffective.

5.1.2 Hedging at the portfolio level

Financial options are usually managed with reference to a hedging portfolio, which is traded to minimise the overall possible loss as no hedge is perfect. The risks can never be totally hedged, but they can be handled efficiently (reducing transaction costs) and effectively (by diversifying away untraded risks) when combined and managed at a portfolio level. If a portfolio approach is taken to owning options then hedging individual options is not required. This may be due to independence of the option, or through risk control at the portfolio level.

The more uncorrelated options there are in a portfolio, the closer the actual return of the portfolio will be to its expected return (in relative terms). Independence may also be achieved by repeating the real option over time. The cumulative value of the series of the same option will allow the good and bad payoffs to mitigate one another, removing the need to hedge.

Within a firm there may already be natural hedges, so that hedging a risk that is already hedged will lead to speculation and increased risk. For example, a firm that sells both ice-creams and umbrellas may find that investing in weather derivatives only increases its cashflow volatility.

5.1.3 Comparing real option strategies

Strategies are often framed in a risk-return space, where the optimal strategy is one with the largest return for a given level of risk. The SDF approach allows the risk and return measures to be combined into a single measure of shareholder value, combining all timescales.

This is achieved by adding an appropriate risk charge to the expected value of each strategy. Section 6 demonstrates this approach with a case study.

Introduction

5.2 Expected discounted shortfall

This paper uses the SDF to create a coherent risk measure that views cashflows over multiple time periods and takes into account correlation with the market. Expected discounted shortfall (EDS) can be seen as a discounted extension of expected shortfall (ES, section 2.2.1), applied to a series of cashflows. Another differentiating characteristic is that absolute levels of cash are defined over many timepoints, rather than a percentile level at a single point in time. An absolute approach is required as a percentile measure does not easily extend to a multi-temporal environment. A practical example of EDS is given in section 6.

5.2.1 Cashflow default point

A company defaults when its incoming cashflow is insufficient to cover its required cashflow expenditure. It is (theoretically) possible to identify, at each time point, the minimum cashflow that allows the company to continue trading. This is defined as the cashflow default point (CDP). The default point for each time period should be calculated at the firm level, so that defaults occur after all relevant income cashflows are taken into consideration. The actual value of the CDP will be dependent on many factors (gearing level, opex requirements etc.) which are time varying, and could also be random variables. To ease calculation and comparison it may be more convenient to use a range of constant CDP levels as a proxy for the actual default level.

An alternative approach is to think of the CDP as the point at which undesirable distress is caused, rather than complete default. For example, this could be the point at which new debt has to be issued at a high interest rate, or it could be a minimum level set by the company directors. If there are business decisions that can be made at the company level to reduce the probability of a default or distress event occurring then these should be explicitly (and dynamically) modelled. One way to deal with default is to require that assets be sold, but rather than force income generating assets to be sold, an extra injection of shareholder equity is modelled. Enough capital is assumed to be raised by shareholders to cover the shortfall, which is defined as

$$S_t = \max\left(CDP_t - x_t, 0\right) \tag{43}$$

where x_t is the net cashflow generated by the company's assets at time t. The cashflow default point CDP_t is the minimum required outgoing cashflow at time t. This is the amount owed by the company to its creditors, e.g. a coupon payment on a bond.

5.2.2 Expected discounted shortfall (EDS)

Expected discounted shortfall (EDS) is defined as

$$EDS = \sum_t E\left[m_t S_t\right] = \sum_t E\left[m_t \max\left(CDP_t - x_t, 0\right)\right] \tag{44}$$

where CDP_t and x_t are defined above and m_t is the SDF.

5.2.3 Conditioning on default

After finding the EDS cost to the company, a project can be charged or rewarded for the default risk it adds or removes. This charge only makes sense from a company perspective, since the entire company, not just the one project, is in default.

Defining x_t^i as the cashflow of the ith business area at time t allows the company cashflow at time t to be written as $x_t = \sum_i x_t^i$. By definition, the risk contributions of individual exposures should add up to the total risk. EDS can be re-expressed as

$$EDS = \sum_t E[m_t S_t] = \sum_t (E[CDP_t m_t | x_t < CDP_t] - E[m_t x_t | x_t < CDP_t])$$

Splitting out the individual components gives $\sum\limits_t E[m_t CDP_t | x_t < CDP_t] - EDS = \sum\limits_i EDS^i$ where $EDS^i = \sum\limits_t E\left[m_t x_t^i | x_t < CDP_t\right]$.

EDS^i can be interpreted as the ith business area's marginal contribution to the EDS of the company. Without the contribution of the ith business area, the company's EDS charge would be EDS^i greater. To calculate the marginal EDS in practice, a Monte Carlo engine is run and all the simulation paths where the company cashflow is always above the CDP are discarded. Using the remaining simulations, the average of the discounted cashflows of each business area can easily be calculated.

One of the main practical difficulties in applying this conditional risk measure is its slow convergence in simulation, especially if the cashflow default point is set at a low level which is not often breached.

5.2.4 EDS capital charge

The EDS charge is the risk the shareholder bears when a project is taken on by a company. A project's default burden to the company can be calculated using the EDS measure, which can be subtracted directly from the project's NPV (see the example in section 6.4). Risks that are negatively correlated with other projects in the company will produce a lower EDS value than positively correlated risks. EDS is an extra cost that should be integrated into the decision making process.

EDS is also dependent on the other projects in a firm's portfolio. One project may have its EDS charge reduced (or even turned positive) when a second complementary project is added to the portfolio. A corollary is that a project that fits well within one company's portfolio is unlikely to find the same fit in a different company.

5.2.5 Properties of EDS

[Csoka et al., 2007] show that any risk measure is coherent if it is based on the basic pricing equation[9] (eq: 4). EDS uses this equation so it is a coherent risk measure. Only the assumptions of free portfolio formation and the law of one price (see section 3.3.1) are required to use EDS.

The SDF process m_t can be used to discount cashflows from all projects and time horizons simultaneously, combining market and non-market risks. Path dependent interventions and outcomes are naturally incorporated. Rather than the risk being controlled by the size of the discount rate, the risk premium is controlled by the covariance of the cashflow with the SDF. Covariance can be set independently for each risk source, allowing different risk premiums to be valued using the same SDF.

EDS is the amount of additional equity guaranteed by the shareholders to keep company ownership from passing to the debt holders. It is the amount that would have to be paid to insure against default at multiple

[9]Referred to in their paper as the general equilibrium measure of risk.

future points, where the insurance covers all default losses of the company. Shareholders must pay out on this insurance to retain ownership of the company. The capital structure of the firm affects the EDS measure. The more debt that is held, the higher the CDP and the greater the EDS charge. The risks of a project may have different effects on EDS at different debt levels.

It may be possible to calculate EDS analytically in some situations, however the real power of EDS lies in its ability to be applied to high dimensional, complex problems. It fits neatly into a Monte Carlo framework allowing complex situations to be modelled. If calculated using Monte Carlo simulation then scenarios that cause distress can be used to identify the risks to a firm. If a recurring risk source can be identified then steps can be taken to mitigate that risk.

EDS can be calculated simultaneously over many default levels and time points, allowing a clear, consistent picture of the risk profile to be formed. Using EDS allows risk and return to be combined into one measure, valued from the shareholder perspective. This allows judgement between various real option strategies, where the optimal strategy is the one generating the greatest shareholder value. EDS does not find optimal strategies, it only provides a way to rank strategies.

5.2.6 Comparison with other approaches

Risk measures that look at a single percentile value ignore any cashflow information occurring outside that percentile. The EDS and ES measures take an average of all the values in the tail, so no information is ignored.

Most other risk measures (ES, VaR) chose a percentile to represent the default point, but EDS defines a (possibly random) default cashflow level, the point at which incoming cashflow is less than the required outgoing cashflow.

A large advantage that EDS has over other measures is that it can be used in a multi-period setting because all shortfalls are discounted back to a present value. Risk and return can be measured in the same units and linked directly to shareholder value.

EDS is mostly relevant at the portfolio level as its most useful economic meaning is when the company is financed by a capital structure that includes defaultable payments. EDS requires a trigger event, as essentially the method is trying to price a form of insurance.

6 Application

The choice between debt and equity involves trading off the benefits of debt (e.g. interest payments paid before tax, increased discipline on managers) against the potential costs of financial distress. To calculate the benefit of debt a firm needs to know the probability of distress occurring, for a given capital structure. One of the most important determinants in calculating the probability of distress is the variability of cashflows.

This section describes an application of the SDF approach to valuing real options, using EDS to measure the risk of the option. The examples in this section have been designed to demonstrate the ideas introduced in this paper. Detailed calculations for parameters are given, allowing Monte Carlo output to be compared to theoretical calculations.

Section 6.1 introduces the situation and section 6.2 derives parameters for Normally distributed mean reverting exponential processes and the correlation parameters for the covariance matrix. Section 6.3 introduces the real option to expand production. Section 6.2.1 reviews the statistics and risk measures. Section 6.5 looks at different possible real options strategies and makes comparisons.

Dynamic risk models have preprogrammed actions that are triggered depending on events occurring. This allows interventions over time to be modelled, according to pre-set strategies. By testing many different combinations of interventions, the best strategy can be identified. Section 6.5.2 uses a dynamic risk model to switch production on or off dependent on two underlying costs. This is a simplified example and in practice the combination of underlying events on which to base decisions is likely to be infinite.

Including embedded optionality into a risk model makes modelling trickier, and optimization even harder but modelling flexibility is essential if solutions, as well as problems, are to be identified. To avoid loss of value it is essential that risk, as well as return, is measured.

6.1 Background information

A widget company produces 5000 widgets every month. All widgets are sold each month at a wholesale price set by an external agent. One of the company's engineers has discovered that production can be increased if part of the process is continually heated to a high temperature using natural gas. At the start of the month the gas supplier fixes the gas price at the market spot rate. Each extra widget requires 0.9 units of gas. The average price of both gas and widgets is 1, so the expected profit from a gas produced widget is 0.1. Management have determined that widget and gas prices both carry a risk premium of 3% and that their price movements are uncorrelated. Interest is assumed to be deterministic with an annual risk free return of $\exp(0.05) = 1.0513$.

6.2 Parameter derivation

The annual variances (measured weekly, multiplied by 52) are

Without Mean Reversion	Widget	Gas
Variance	0.12908	1.88224
Mean reversion $= \alpha$	0.5	4.8
$StdDev[S] = \sqrt{E[S^2] - E[S]^2}$	27.3366%	46.9172%
Mean Reverting LogNormal σ	33.7654%	138.2099%

The calculation for σ is achieved by inverting the equation for the standard deviation of a mean reverting log normal distribution $StdDev[S] = E[S]\sqrt{(\exp(\widehat{\sigma}^2) - 1)}$ where $\widehat{\sigma}^2 = \frac{\sigma^2}{2\alpha}(1 - \exp(-2\alpha))$, giving $\sigma = \sqrt{\ln\left(\left(\frac{StdDev[S]}{E[S]}\right)^2 + 1\right)\frac{2\alpha}{(1 - \exp(-2\alpha))}}$

The minimum variance and correlations of the stochastic discount factor can now be calculated using eq: 30.

	α	MR LogNormal σ	R_i^e	Sharpe Ratio
Widget	0.5	33.77%	3%	0.1118
Gas	4.8	138.20%	3%	0.0673

where $Sharpe\ Ratio_i = \frac{\ln(1+R_i^e)}{\sigma_i}$. Choosing $Maximum\ Correlation = 0.636$ gives $SDF\ Volatility = \sigma_m = \frac{\max(Sharpe\ Ratio_i)}{Maximum\ Correlation} = \frac{0.1118}{0.636} = 13.76\%$ and using $SDF\ Correlation_i = -\frac{Sharpe\ Ratio_i}{SDF\ Volatility}$ produces Table 1

The excess market portfolio return can be calculated as $R_p^e = \exp\left(\sigma_m^2\right) - 1 = 1.9\%$. It should be noted that this is an arbitrary number as it is related to the $Maximum\ Correlation$ parameter. Assuming a higher market return implies a lower correlation and vice-versa.

Table 1: Instantaneous Correlation Matrix

Correlation	SDF	Widget	Gas
SDF	1	-0.6360	-0.1554
Widget	-0.6360	1	0
Gas	-0.1554	0	1

6.2.1 Distribution of widget prices

The Monte Carlo output of widget prices has the statistical properties given in Table 2.

Table 2: Spread of Widget Prices

Widget Prices	Year 1	Year 2	Year 3	Year 4	Year 5
Mean	0.9985	0.9986	1.0020	1.0009	0.9984
St Dev	0.2722	0.3216	0.3393	0.3459	0.3456
5%	0.6175	0.5650	0.5487	0.5445	0.5425
10%	0.6817	0.6352	0.6193	0.6146	0.6127
50%	0.9638	0.9504	0.9510	0.9453	0.9426
90%	1.3591	1.4193	1.4466	1.4544	1.4494
95%	1.4958	1.5925	1.6318	1.6373	1.6416

The percentile measures can be read as the probability that the widget price is below the value given in the table. The spread of widget prices stabilises as time progresses.

6.2.2 Comparison to deterministic NPV

Cash is assumed to be collected at the end of the month and deposited in the bank, so at the end of the year the mean amount collected will be

$$5,000 \sum_{t=0}^{11} \exp\left(\left(11/12 - t\right) 0.05\right) = 5,116.44$$

The market betas in Table 3 for the widget and gas price are calculated using eq: 46 in Appendix A.4. The deterministic discount factors are calculated using CAPM as $D_t^i = \exp\left(-\left(\beta_i r_p^e + r_f\right) t\right)$ where $r_p^e = \ln\left(R_p^e + 1\right)$.

Table 3: Beta and Discount factors

Time $= t$	0	1/12	2/12	3/12	4/12	5/12	6/12	...
Widget β_w	1.560	1.528	1.497	1.467	1.437	1.409	1.381	...
Gas β_g	1.560	1.286	1.074	0.909	0.778	0.675	0.591	...
Widget D^w	1.000	0.993	0.987	0.981	0.975	0.969	0.963	...
Gas D^g	1.000	0.994	0.988	0.983	0.979	0.974	0.970	...

Table 4: Monte Carlo Mean Values

End of year means	Year 1	Year 2	Year 3	Year 4	Year 5	s.e. Year 1
SDF	0.9517	0.9061	0.8619	0.8202	0.7811	0.0006
Widget Price	0.9980	0.9986	1.0022	1.0012	0.9990	0.0012
Gas Price	0.9997	0.9995	0.9996	0.9977	1.0030	0.0021
Production	5,000	5,000	5,000	5,000	5,000	0
Cash	4,998	4,995	5,000	5,006	4,999	3.4821
Cash + Interest	5,115	5,111	5,116	5,123	5,116	3.5402
Disc Cash	4,830	4,509	4,245	4,017	3,802	2.6056

Assuming that each month has a constant $1/12$ of annual production then the NPV for the 1st year of widget sales can be calculated, using the monthly discount factors, as $5000 \sum_{t=0}^{11} D^w_{t/12} = 4831.35$ and further years as:-

Theoretical	Year 1	Year 2	Year 3	Year 4	Year 5
NPV by year	4,831	4,511	4,242	4,007	3,796

Performing $50,000$ Monte Carlo simulations gives the results in table 4. The line *SDF* is the mean value of the stochastic discount factor, where the average is designed to converge to the risk free rate. The final column of the table shows the standard error of Year 1 simulated values. As expected, *Gas* and *Widget* prices are all within 2 standard errors (s.e.). Production is not a random variable so has zero standard error. *Cash* is the total cash collected in each year, without discounting or interest. *Cash + Interest* is the total cash at the end of each year after being deposited at the risk free rate at the end of each month. *Disc Cash* is the present value of the cashflows received discounted using the stochastic discount factor, all of which are within 2 s.e. of the theoretical amounts. These results show that Monte Carlo methods produce reasonable results in comparison to the theoretical calculations. The real options described in the rest of this section are measured against these base case values.

6.3 Valuing the option to expand

This section details the theoretical value gained by increasing widget production using the new gas technology but ignoring any increase in risk. The theoretical NPV when production is increased can be calculated using

$$NPV = 5000 \sum_{i=1}^{60} \left((1 + Production\ Inc.)\, D^w - 0.9 * Production\ Inc. * D^G \right) /12$$

to give the expected present value when using the option, ignoring risk factors. The theoretical calculations are presented in Table 5 with the Monte Carlo outputs presented in Table 6.

Valuation of mean reverting cashflows using the SDF is much simpler to implement than theoretical calculations, so long as the problem is large and complex enough to warrant a Monte Carlo approach.

On average, creating an extra widget is NPV positive so if risk is ignored then value is maximised by increasing production to maximum capacity. The risk dimension should not be ignored as the next section will show.

Table 5: Theoretical NPV

Production Inc.	NPV Inc.	NPV increase broken down by year				
		Year 1	Year 2	Year 3	Year 4	Year 5
10%	155	45	35	29	24	22
30%	466	136	106	86	73	65
50%	777	226	176	143	122	108
70%	1,087	317	246	201	171	152
90%	1,398	408	317	258	220	195

Table 6: Monte Carlo NPV

Production Increase	NPV Increase	NPV increase broken down by year				
		Year 1	Year 2	Year 3	Year 4	Year 5
10%	156	45	34	29	25	22
30%	468	136	103	87	76	65
50%	779	227	172	146	127	108
70%	1,091	318	241	204	177	151
90%	1,403	409	310	262	228	194

6.4 Calculating the cost of risk

If the owners of the company had unlimited access to funds then risk would not be an issue. All decisions could be made on an expected NPV basis. In reality this is rarely the case. In this example case study the project has been implemented by a financial structure that requires a fixed annual payment[10]. Calculations are given at various payment levels to demonstrate different risk profiles.

Table 7: Expected Shortfall, risk free rate

Shortfall Level	Year 1	Year 2	Year 3	Year 4	Year 5
3,800	5	75	167	**264**	359
4,000	13	116	247	**382**	514
4,200	29	179	354	**536**	713
4,400	57	258	494	**731**	962
4,600	100	366	**669**	971	1,623
4,800	161	504	**884**	1,260	1,623
5,000	242	674	**1,140**	1,598	2,040

Table 7 shows the expected shortfall at various levels, discounted to the present at the deterministic risk free rate and where no extra widgets are produced. The values in the table represent the cumulative mean amount of cash that would need to be set aside now to cover the expected value of a possible shortfall in cash.

[10] For example, this could be a coupon payment of a bond or a project financing requirement

As the shortfall level increases so does the cumulative expected shortfall because the volatility of the widget price makes the cashflow more likely to fall below this level. The expected shortfall for a specific year can be calculated by subtracting one cumulative value from another. For example, the cost at shortfall level 4400 at time 4 is $731 - 494 = 237$. The **bold** values indicate the year when the greatest increase in premium occurs. This point is the riskiest point in time, for the given shortfall level. Discounting the expected shortfall at the risk free rate may seem like the natural, prudent approach but this ignores the risk of losses occurring at a time when the rest of the market is also falling.

Table 8 displays the Expected Discounted Shortfall with discounting using the stochastic discount factor rather than the risk free rate. The values represent the 'insurance premiums' payable at time 0 for protection below the shortfall level over the period 1-5 years, where if the annual cashflow is lower than the shortfall level then the EDS insurance makes up the difference. The insurance can be invoked in multiple years, i.e. if it pays in year 2 it can also pay in year 4.

Table 8: EDS over 5 years

Shortfall Level	Year 1	Year 2	Year 3	Year 4	Year 5
$3,800$	6	92	211	**338**	464
$4,000$	16	141	306	**480**	652
$4,200$	34	210	431	**661**	888
$4,400$	65	303	591	**886**	$1,177$
$4,600$	113	424	789	**$1,159$**	$1,523$
$4,800$	179	576	$1,027$	**$1,482$**	$1,927$
$5,000$	265	759	**$1,307$**	$1,854$	$2,389$

The premiums demanded by the EDS measure are all greater than the ES measure discounted at the risk free rate. This implies that discounting at the risk free rate will underestimate risk for a diversified investor. In fact, if a deterministic discount rate is used to match the risk free discounted values to the EDS values then the rate would be around 2% less than the risk free rate.

The cost of insurance is most expensive in the 4th year as this is the most risky year for all shortfall levels except $5,000$. For shortfall levels $4,600$ and $4,800$ the riskiest year has changed. So not only is the size of risk incorrect when using the risk free rate but also the point of maximum risk is incorrect. Decisions made after discounting risks at the risk free rate can lead to suboptimal strategies.

6.5 Options as strategies

6.5.1 Choosing between strategies

As well as comparing risks over time, EDS allows different real option strategies to be compared consistently. Table 9 compares using gas for increased production to the base case of no extra production[11]. The **bold** values are the years when the extra risk introduced by increased production is greater than the standalone NPV of the extra production. The cost of shortfall insurance for producing these extra widgets is greater than the value they add.

At all production levels the EDS premium for year 4 is the greatest. In year 5 present value of the risk is lower than in year 4 even though the probability of the cashflow going below the shortfall level is larger. The

[11] The comparison base case for this table is with 0% extra production.

Table 9: EDS

Shortfall Level			5000	4800	4600	4400	4200	4000	3800
Prop	Year	Inc NPV	Increase in EDS compared to no gas						
10%	1	45	9	12	13	12	10	7	5
10%	2	34	30	32	32	32	30	28	24
10%	3	29	**35**	**36**	**37**	**36**	**35**	**33**	**30**
10%	4	25	**36**	**38**	**38**	**38**	**36**	**34**	**31**
10%	5	22	**35**	**36**	**37**	**36**	**35**	**33**	**31**
30%	1	136	33	37	39	36	31	25	19
30%	2	103	99	**103**	**105**	**104**	100	93	84
30%	3	87	**112**	**116**	**118**	**117**	**114**	**108**	**101**
30%	4	76	**116**	**120**	**121**	**120**	**117**	**112**	**105**
30%	5	65	**112**	**115**	**117**	**116**	**113**	**108**	**101**
50%	1	227	45	48	48	46	42	37	30
50%	2	172	**176**	**183**	**186**	**184**	**179**	169	156
50%	3	146	**196**	**202**	**205**	**205**	**200**	**192**	**181**
50%	4	127	**202**	**207**	**210**	**209**	**205**	**198**	**187**
50%	5	108	**194**	**200**	**202**	**202**	**198**	**191**	**181**
70%	1	318	51	53	53	52	49	44	39
70%	2	241	**258**	**267**	**271**	**269**	**262**	**250**	234
70%	3	204	**284**	**292**	**296**	**296**	**291**	**281**	**267**
70%	4	177	**291**	**299**	**302**	**301**	**297**	**288**	**275**
70%	5	151	**281**	**288**	**291**	**291**	**286**	**278**	**266**
90%	1	409	55	56	56	55	53	49	45
90%	2	310	**344**	**354**	**359**	**357**	**349**	**353**	**316**
90%	3	262	**376**	**385**	**390**	**390**	**384**	**373**	**358**
90%	4	228	**384**	**392**	**396**	**396**	**391**	**381**	**366**
90%	5	194	**370**	**378**	**382**	**382**	**377**	**367**	**353**

more that production is increased, the greater the NPV, but also the larger the risk. The cost of risk increases faster than the NPV. In the long term, continually using gas to increase production is not an optimal strategy.

6.5.2 Option to switch on and off

The engineer has modified the process so that the gas does not need to be continually switched on. To get the equipment up to temperature after an inactive period requires 1.5 units of gas per widget for the initial month, reverting back to 0.9 units per widget for subsequent months. Extra production can be shut down at no cost, but restarting requires another month using the extra 1.5 units of gas.

The new strategy, tested in Table 10, is to switch on extra gas production when widget prices are 0.5 higher than gas prices, and switch off the extra production when widget prices are less than gas prices. This strategy is guaranteed to either equal or better the base case of no extra production.

Using the new strategy to choose when to turn switch on, increases the value of the project. Not only does the NPV increase but the EDS premium also decreases, making the whole project less risky. This is not surprising as the strategy is designed never to make a loss (in comparison to the base case). The best option is to maximise production at 90% when gas is 0.5 cheaper than widget sales prices.

Table 10: New Option

Shortfall Level			5000	4800	4600	4400	4200	4000	3800
Prop	Year	Inc NPV	Increase in EDS compared to no gas						
Normally distributed random variables									
90%	1	584	−9	−6	−3	−2	−1	0	0
90%	3	762	−97	−75	−56	−40	−27	−18	−11
90%	4	738	−89	−71	−54	−40	−28	−19	−12
90%	5	701	−82	−65	−50	−37	−26	−18	−11
90%	5	655	−79	−63	−49	−36	−26	−18	−11

Being equal or better than the base case is not necessarily the optimal strategy. It may be better to start extra production up when gas prices are only, say, 0.45 cheaper than widget prices so a small loss is made in the first month allowing the 0.9 production to be ready for subsequent months. There may also be options available to reduce (or increase) the required payment in the financing framework. Meaningful decisions can only be made once the risk profile changes have had their costs and benefits analyzed.

6.5.3 Hedging Strategies

The sale and purchace of relevant financial products can be added alongside the above analysis. This allows the effect of hedging strategies to be measured at the default level, rather than at the option level. The effect of different strategies can be measured and compared, allowing partial hedges to be evaluated and the costs of a particular strategy to be compared to the EDS saving.

7 Conclusion

7.1 Stochastic Discount Factors

The stochastic evolution of asset prices has become accepted as a mainstream valuation and risk measurement technique within corporate finance, so there is no reason why a stochastic discount rate should not be equally as accepted, especially when the discount rate is strongly linked to the accepted concept of the market portfolio.

The value of an SDF is related to the performance of the market, not to the risks of a project. By assigning a high discount rate when the market is high and a low discount rate when the market is low, the SDF rewards cashflows received when the economy is falling and punishes cashflows received when the economy is rising. This is the shareholders' view of risk, as a risk-averse investor prefers a smooth series of cashflows to a volatile series and so will be willing pay more for an asset that will increase in value when the majority of other investments are decreasing in value.

SDFs are mathematically equivalent to risk neutral, certainty equivalent and CAPM approaches. Conceptually, correlation to a market portfolio is much easier to understand than alternative justifications, allowing greater chance of adoption by practitioners. SDFs use real world probabilities to project realistic growths in asset value, so are easier to communicate. SDFs also ensure that models are arbitrage free.

The SDF models described in the literature are either too specific or too complex to implement, so have limited the appeal of the SDF method. This paper showed that a continuous time, multivariate model can be built where the mathematical complexities of other SDF methods are replaced with a covariance definition, giving a simple but strong foundation for the development of more advanced SDF models. This was used to define a lognormal SDF model with formulae for controlling the risk premia through the correlation matrix. This simple implementation is important for practical adoption of the SDF method, outside of empirical research. The multivariate covariance pricing model of section 3.5.1 placed only a few restrictions on the price and SDF processes. It made no assumptions about probability distributions, although it implicitly assumed that the mean, variance and correlation could be calculated.

Equations 25 and 30 are important as they explicitly link the risk premium to the correlation parameter of the SDF and the asset.

Monte Carlo SDF modelling does not provide quick mathematical solutions or define a replicating portfolio, but neither of these are usually essential for real options evaluation.

7.2 Real options

Using the SDF technique within a Monte Carlo framework allows real options to be valued using real world probabilities, without having to identify a replicating portfolio. Equilibrium based economic arguments can be applied instead, avoiding all difficulties surrounding the tradeability and hedging of an asset. SDFs are dependent on the choice of statistical model, but are independent of the cashflows being valued. This independence allows the same SDF model simulations to be applied to multiple projects, reducing the number of models that have to be communicated to management and allowing the same SDF set to be used for many different real options at company, project and even transaction level. SDFs should also be used when cashflows are truncated by guarantees or contractual agreements, or when the shareholders risk point of view requires consideration.

Approaches that combine traded and untraded risks usually sacrifice some of the specific information associated with each of these types of uncertainty. By removing the hedging requirement and using equilibrium arguments, the discount rate implied by the SDF is set in relation to the asset's correlation with the market, allowing traded and untraded risks to be treated in the same manner.

SDFs remove the complexity and subjectivity around choosing a discount rate for project valuation. Fixing the risk premiums for specific risks allows a standard approach to be applied to all projects. These risk premiums can be calibrated either by an assumed risk premium or by historical correlation with a proxy for the market portfolio. These assumptions may be decided centrally, allowing practitioners to concentrate on modelling options.

Unfortunately the risk of a real option is not controlled under the Monte Carlo SDF approach. This has to be handled through other methods, such as the expected discounted shortfall risk measure.

7.3 Expected discounted shortfall

Risk measures in the current literature do not allow risks in different time periods to be compared consistently. Most risk measures treat risks that are correlated with the market exactly the same as risks that are uncorrelated with the market, therefore they do not take into account shareholder risk preferences. Ultimately, for a company to add value, the shareholder perspective needs to be the overriding measure, taking in risk of default as well as correlation with the market. Expected discounted shortfall (EDS) is a new coherent risk measure that is able to combine multiple time periods to measure risk from the point of a diversified shareholder. Risk is measured in terms of value, not probability, making it ideal to judge the effectiveness of a particular strategy, and compare between strategies.

Using SDFs in combination with a cashflow model allows risk reduction to be judged across strategies at the company level, whilst also quantifying any value created by the risk reduction. Risk-return decisions are greatly simplified as the risk dimension has been converted into value. This paper showed that the risk added by a project can outweigh the expected value added.

Discounting the expected shortfall at the risk free rate ignores the risk of losses occurring at a time when the rest of the market is also falling, so underestimates the risk to a diversified shareholder. As the discount rate is incorrect, so are estimates of the point of maximum risk. This paper shows that this can lead to resources being mis-allocated.

There are a number of equivalent ways to interpret EDS:-

- It is the present value of the average amount the equity shareholders may have to inject into the company to keep it as a going concern.
- If the company were able to buy insurance covering all defaults for a given period then the EDS would be the fair value for a one-off premium.
- If default payments are completely hedgeable, then the EDS equals the present value of an (unidentified) replicating hedge portfolio.

The EDS approach to real options does not provide the optimal solutions, but it does give a consistent framework where different option strategies can be compared, combining present values of expected payoffs with the present value of their associated risk costs.

7.4 Paper Innovations

This paper employs generic SDF modelling framework, based around covariance to derive a log-normally distributed, mean reverting, multi-asset model that gives an analytic relationship between the risk premium of each asset and the SDF, using the correlation matrix.

Applying the SDF to corporate finance simplifies many problems, particularly when used with real options and risk calculations. The Monte Carlo SDF approach is simple for real option practitioners to implement and explain as it uses real world probabilities. The approach uses equilibrium theory, avoiding assumptions that have hampered the acceptance of real options.

Expected discounted shortfall is a new, coherent, multiperiod risk measure that is applied at the company level, but can be measured at the transaction level. It is the only risk measure that is able to coherently combine risks from multiple time periods into a single present value, making it the ideal measure to judge the effectiveness of different real option strategies, from the shareholders point of view. The ability to convert risk into present value allows risk and return to be compared on the same basis over multiple time periods.

7.5 Further research - Optimality

Optimal real options strategies are not identified using the Monte Carlo method. Applying stochastic programming alongside SDF method would allow shareholder value to be maximised and adding the element of competition to the SDF approach would enhance decision making abilities. Factor modelling can reduce the number of random variables and increase efficiency. Methods other than Monte Carlo need to be employed for optimal strategies to be identified.

References

[Acerbi & Tasche, 2002] Acerbi, C. & Tasche, D. (2002). Expected Shortfall: A Natural Coherent Alternative to Value at Risk. *Economic Notes*, 31(2), 379–388. arXiv:cond-mat/0105191v1.

[Ait-Sahalia & Lo, 1998] Ait-Sahalia, Y. & Lo, A. W. (1998). Nonparametric estimation of state-price densities implicit in financial asset prices. *The Journal of Finance*, 53(2), 499–547. http://www.jstor.org/stable/117360.

[Alvarez & Jermann, 2005] Alvarez, F. & Jermann, U. J. (2005). Using asset prices to measure the persistence of the marginal utility of wealth. *Econometrica*, 73(6), 1977–2016. http://www.jstor.org/stable/3598756.

[Amram & Kulatilaka, 1999] Amram, M. & Kulatilaka, N. (1999). *Real Options: Managing Strategic Investment in an Uncertain World*. Boston, Massachusetts: Harvard Business School Press. ISBN-10: 0875848451.

[Arajo et al., 2005] Arajo, F., Issler, J. V., & Fernandes, M. (2005). *Estimating the Stochastic Discount Factor without a Utility Function*. Technical Report 583, Graduate School of Economics, Getulio Vargas Foundation. http://ideas.repec.org/p/fgv/epgewp/583.html.

[Arrow & Debreu, 1954] Arrow, K. J. & Debreu, G. (1954). Existence of an Equilibrium for a Competitive Economy. *Econometrica*, 22(3), 265–290. http://www.jstor.org/stable/1907353.

[Artzner et al., 1999] Artzner, P., Delbaen, F., Eber, J.-M., & Heath, D. (1999). Coherent Measures of Risk. *Mathematical Finance*, 9(3), 203–228. doi:10.1111/1467-9965.00068.

[At-Sahalia & Lo, 2000] At-Sahalia, Y. & Lo, A. W. (2000). Nonparametric risk management and implied risk aversion. *Journal of Econometrics*, 94(1-2), 9–51. doi:10.1016/S0304-4076(99)00016-0.

[Bakshi et al., 2008] Bakshi, G., Carr, P., & Wu, L. (2008). Stochastic risk premiums, stochastic skewness in currency options, and stochastic discount factors in international economies. *Journal of Financial Economics*, 87(1), 132–156. doi:10.1016/j.jfineco.2006.12.001.

[Bekaert & Grenadier, 2002] Bekaert, G. & Grenadier, S. R. (2002). Stock and Bond Pricing in an Affine Economy. *EFA 2002 Berlin Meetings*. http://ssrn.com/paper=299952.

[Bernardo & Ledoit, 2000] Bernardo, A. E. & Ledoit, O. (2000). Gain, loss and asset pricing. *The Journal of Political Economy*, 108(1), 144–172. http://www.jstor.org/stable/3038246.

[Black & Scholes, 1973] Black, F. & Scholes, M. (1973). The pricing of options and corporate liabilities. *The Journal of Political Economy*, 81(3), 637–654. http://www.jstor.org/stable/1831029.

[Borison, 2005] Borison, A. (2005). Real Options Analysis: Where Are the Emperor's Clothes? *Journal of Applied Corporate Finance*, 17(2), 1731. doi:10.1111/j.1745-6622.2005.00029.x.

[Breymann et al., 2005] Breymann, W., Kelly, L., & Platen, E. (2005). Intraday Empirical Analysis and Modeling of Diversified World Stock Indices. *Asia-Pacific Financial Markets*, 12(1), 1–28. doi:10.1007/s10690-006-9010-0.

[Campbell, 2003] Campbell, J. Y. (2003). Consumption-based asset pricing. *Handbook of the Economics of Finance*, 1(2). doi:10.1016/S1574-0102(03)01022-7.

[Chapman, 1997] Chapman, D. A. (1997). Approximating the Asset Pricing Kernel. *The Journal of Finance*, 52(4), 1383–1410. http://www.jstor.org/stable/2329440.

[Chapman et al., 2001] Chapman, R., Gordon, T. J., & Speed, C. A. (2001). Pensions Funding and Risk. *British Actuarial Journal*, 7(4), 605–662. www.ingentaconnect.com/content/fia/baj/2001/00000007/00000004/605_662.

[Clewlow & Strickland, 1999] Clewlow, L. & Strickland, C. (1999). *Valuing Energy Options in a One Factor Model Fitted to Forward Prices*. Research Paper Series 10, Quantitative Finance Research Centre, University of Technology, Sydney. http://ideas.repec.org/p/uts/rpaper/10.html.

[Clewlow & Strickland, 2000] Clewlow, L. & Strickland, C. (2000). *Energy derivatives - Pricing and risk management*. Australia: Lacima Group. ISBN-10: 0953889602.

[Cochrane, 2004] Cochrane, J. H. (2004). *Asset Pricing*. New Jersey: Princeton University Press. ISBN-10: 0691121370.

[Cochrane & Sa-Requejo, 2000] Cochrane, J. H. & Sa-Requejo, J. (2000). Beyond arbitrage: Good-deal asset price bounds in incomplete markets. *The Journal of Political Economy*, 108(1), 79–119. http://www.jstor.org/stable/3038244.

[Constantinides, 1992] Constantinides, G. M. (1992). A theory of the nominal term structure of interest rates. *The Review of Financial Studies*, 5(4), 531–552. http://www.jstor.org/stable/2962140.

[Copeland & Antikarov, 2001] Copeland, T. E. & Antikarov, V. (2001). *Real Options: A Practitioner's Guide*. New York: W. W. Norton & Company, first edition. ISBN-10: 1587990288.

[Csoka et al., 2007] Csoka, P., Herings, J.-J. P., & Koczy, L. . (2007). Coherent Measures of Risk from a General Equilibrium Perspective. *Journal of Banking & Finance*, 31(8), 2517–2534. http://mek.oszk.hu/05300/05310/05310.pdf.

[Dai & Singleton, 2000] Dai, Q. & Singleton, K. J. (October 2000). Specification analysis of affine term structure models. *The Journal of Finance*, 55, 1943–1978. doi:10.1111/0022-1082.00278.

[Dixit & Pindyck, 1994] Dixit, A. K. & Pindyck, R. S. (1994). *Investment under Uncertainty*. New Jersey: Princeton University Press. ISBN-10: 0691034109.

[Duffie, 2001] Duffie, D. (2001). *Dynamic Asset Pricing Theory*. New Jersey: Princeton University Press, third edition. ISBN-10: 069109022X.

[Duffie et al., 2000] Duffie, D., Pan, J., & Singleton, K. (2000). Transform Analysis and Asset Pricing for Affine Jump-Diffusions. *Econometrica*, 68(6), 1343–1376. http://www.jstor.org/stable/3003992.

[Epstein & Zin, 1991] Epstein, L. G. & Zin, S. E. (1991). Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis. *Journal of Political Economy*, 99(2), 263–86. http://www.jstor.org/stable/2937681.

[Etheridge, 2002] Etheridge, A. (2002). *A course in Financial Calculus*. Cambridge: Cambridge University Press. ISBN 0521890772.

[Ferson, 1995] Ferson, W. E. (1995). Theory and Empirical Testing of Asset Pricing Models. *Handbooks in Operations Research and Management Science*, 9, 145–200. doi:10.1016/S0927-0507(05)80049-0.

[Fisher, 1907] Fisher, I. (1907). *Rate of Interest: Its nature, determination and relation to economic phenomena*. New York: Macmillan.

[Froot & Stein, 1998] Froot, K. & Stein, J. (1998). Risk management, capital budgeting, and capital structure policy for financial institutions: An integrated approach. *Journal of Financial Economics*, 47(1), 55–82. doi:10.1016/S0304-405X(97)00037-8.

[Girotto & Ortu, 1997] Girotto, B. & Ortu, F. (1997). Numeraires, equivalent martingale measures and completeness in finite dimensional securities markets. *Journal of Mathematical Economics*, 27(3), 283–294. doi:10.1016/S0304-4068(96)00779-3.

[Hadzi-Vaskov & Kool, 2008] Hadzi-Vaskov, M. & Kool, C. J. (2008). *Stochastic Discount Factor Approach to International Risk-Sharing: Evidence from Fixed Exchange Rate Episodes*. Working Papers 07-33, Utrecht School of Economics. http://ideas.repec.org/p/use/tkiwps/0733.html.

[Hansen & Jagannathan, 1991] Hansen, L. P. & Jagannathan, R. (1991). Implications of security market data for models of dynamic economies. *The Journal of Political Economy*, 99(2), 225–262. http://www.jstor.org/stable/2937680.

[Hansen & Jagannathan, 1997] Hansen, L. P. & Jagannathan, R. (1997). Assessing Specification Errors in Stochastic Discount Factor Models. *The Journal of Finance*, 52(2), 557–590. http://www.jstor.org/stable/2329490.

[Harrison & Pliska, 1981] Harrison, J. M. & Pliska, S. R. (1981). *Martingales and Stochastic Integrals in the Theory of Continous Trading*. Discussion Papers 454, Northwestern University, Center for Mathematical Studies in Economics and Management Science. http://ideas.repec.org/p/nwu/cmsems/454.html.

[Henderson, 2007] Henderson, V. (2007). Valuing the option to invest in an incomplete market. *Mathematics and Financial Economics [1862-9679]*, 1(2), 103 –128. http://ssrn.com/paper=569865.

[Henderson & Hobson, 2002] Henderson, V. & Hobson, D. G. (2002). Real options with constant relative risk aversion. *Journal of Economic Dynamics and Control*, 27(2), 329–355. doi:10.1016/S0165-1889(01)00052-5.

[Ho & Liu, 2002] Ho, S. P. & Liu, L. Y. (2002). An option pricing-based model for evaluating the financial viability of privatized infrastructure projects. *Construction Management and Economics*, 20, 143–156(14). doi:10.1080/01446190110110533.

[Hodder et al., 2001] Hodder, J. E., Mello, A. S., & Sick, G. (2001). Valuing Real Options: Can Risk-Adjusted Discounting Be Made to Work. *Journal of Applied Corporate Finance*, 14(2), 90101. doi:10.1111/j.1745-6622.2001.tb00333.x.

[Jaimungal et al., 2009] Jaimungal, S., Souza, M., & Zubelli, J. (2009). Real option pricing with mean-reverting investment and project value.

[Jarvis et al., 2001] Jarvis, S., Southall, F., & Varnell, E. (2001). Modern Valuation Techniques. *Staple Inn Actuarial Society*. http://www.sias.org.uk/data/papers/ModernValuations/DownloadPDF.

[Mathews et al., 2007] Mathews, S., Datar, V., & Johnson, B. (2007). A Practical Method for Valuing Real Options: The Boeing Approach. *Journal of Applied Corporate Finance*, 19(2), 95104. doi:10.1111/j.1745-6622.2007.00140.x.

[Mello & Pyo, 2002] Mello, A. S. & Pyo, U. (2002). Real Options with Market Risks and Private Risks. *Journal of Applied Corporate Finance*, 15(2), 89101. doi:10.1111/j.1745-6622.2002.tb00699.x.

[Metcalf & Hassett, 1995] Metcalf, G. E. & Hassett, K. A. (1995). Investment under alternative return assumptions comparing random walks and mean reversion. *Journal of Economic Dynamics and Control*, 19(8), 1471 – 1488. doi:10.1016/0165-1889(94)00838-9.

[Miao & Wang, 2005] Miao, J. & Wang, N. (2005). Investment, Consumption and Hedging under Incomplete Markets. *Boston University, Boston and Columbia Business School, New York*. http://repec.org/sed2006/up.12751.1139632486.pdf.

[Miller & Upton, 1985] Miller, M. H. & Upton, C. W. (1985). A Test of the Hotelling Valuation Principle. *The Journal of Political Economy*, 93(1), 1–25. http://www.jstor.org/stable/1830498.

[Mo & Wu, 2007] Mo, H. & Wu, L. (2007). International capital asset pricing: Evidence from options. *Journal of Empirical Finance*, 14(4), 465–498. doi:10.1016/j.jempfin.2006.12.002.

[Monfort & Gourieroux, 2007] Monfort, A. & Gourieroux, C. S. (2007). Econometric specification of stochastic discount factor models. *Journal of Econometrics*, 136(2), 509–530. doi:10.1016/j.jeconom.2005.11.015.

[Myers, 1976] Myers, S. C. (1976). Using simulation for risk analysis. *Modern Developments in Financial Management*. ISBN-10: 027589360X.

[Ng & Varnell, 2003] Ng, H. M. & Varnell, E. (2003). Frictional costs. *Staple Inn Actuarial Society*. http://www.sias.org.uk/data/papers/FrictionalCosts/DownloadPDF.

[Piazzesi, 2009] Piazzesi, M. (2009). Affine term structure models. *HANDBOOK OF FINANCIAL ECONOMETRICS*, 1(1), Chapter 13. ISBN 044450897X.

[Platen, 2000] Platen, E. (2000). Risk Premia and Financial Modeling without Measure Transformation. http://www.qfrc.uts.edu.au/research/research_papers/rp45.pdf.

[Platen, 2006] Platen, E. (2006). A Benchmark Approach To Finance. *Mathematical Finance*, 16(1), 131–151. doi:10.1111/j.1467-9965.2006.00265.x.

[Robichek & Myers, 1965] Robichek, A. & Myers, S. (1965). *Optimal Financing Decisions*. Englewood Cliffs - New Jersey: Prentice-Hall.

[Rogers, 1997] Rogers, L. C. G. (1997). The Potential Approach to the Term Structure of Interest Rates and Foreign Exchange Rates. *Mathematical Finance*, 7(2), 157–176. doi:10.1111/1467-9965.00029.

[Rosenberg & Engle, 2002] Rosenberg, J. V. & Engle, R. F. (2002). Empirical pricing kernels. *Journal of Financial Economics*, 64(3), 341–372. doi:10.1016/S0304-405X(02)00128-9.

[Rubinstein, 1976] Rubinstein, M. (1976). The Valuation of Uncertain Income Streams and the Pricing of Options. *The Bell Journal of Economics*, 7(2), 407–425. http://www.jstor.org/stable/3003264.

[Sarkar, 2003] Sarkar, S. (2003). The effect of mean reversion on investment under uncertainty. *Journal of Economic Dynamics and Control*, 28(2), 377 – 396. doi:10.1016/0165-1889(94)00838-9.

[Sharpe, 1964] Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *The Journal of Finance*, 19(3), 425–442. http://www.jstor.org/stable/2977928.

[Sick, 1995] Sick, G. (1995). *Real options - Chapter 21*. Elsevier: Handbooks in Operations Research and Management Science. doi:10.1016/S0927-0507(05)80065-9.

[Smith & Speed, 1998] Smith, A. D. & Speed, C. A. (1998). Gauge Transforms. *8th AFIR Colloquium*, (pp. 445–487). http://actuaries.org/AFIR/colloquia/Cambridge/Smith_Speed.pdf.

[Smith & Wickens, 2002] Smith, P. N. & Wickens, M. R. (2002). *Asset Pricing with Observable Stochastic Discount Factors*. Discussion Papers 02/03, Department of Economics, University of York. http://ideas.repec.org/p/yor/yorken/02-03.html.

[Staum, 2004] Staum, J. (2004). Fundamental theorems of asset pricing for good deal bounds. *Mathematical Finance*, 14(2), 141–161. doi:10.1111/j.0960-1627.2004.00186.x.

[Swinand et al., 2005] Swinand, G. P., Rufin, C., & Sharma, C. (2005). Valuing Assets Using Real Options: An Application to Deregulated Electricity Markets. *Journal of Applied Corporate Finance*, 17(2), 5567. doi:10.1111/j.1745-6622.2005.00032.x.

[Theil, 1957] Theil, H. (1957). A Note on Certainty Equivalence in Dynamic Planning. *Econometrica*, 25(2), 346–349. http://www.jstor.org/stable/1910260.

[Thijssen, 2002] Thijssen, J. (2002). A stochastic discount factor approach to investment under uncertainty. http://www.tcd.ie/Economics/TEP/2006_papers/TEP2.pdf.

[Trigeorgis, 1996] Trigeorgis, L. (1996). *Real Options: Managerial Flexibility and Strategy in Resource*. Massachusetts: MIT Press. ISBN-10: 026220102X.

[University of Maryland, 2002] University of Maryland (2002). Roundtable on Real Options and Corporate Practice. *Journal of Applied Corporate Finance*, 15(2), 823. doi:10.1111/j.1745-6622.2002.tb00691.x.

[Varian, 1987] Varian, H. R. (1987). The arbitrage principle in financial economics. *The Journal of Economic Perspectives*, 1(2), 55–72. http://www.jstor.org/stable/194291.

[Yang & Ewald, 2007] Yang, Z. & Ewald, C.-O. (2007). Utility Based Pricing and Exercising of Real Options Under Geometric Mean Reversion and Risk Aversion Toward Idiosyncratic Risk. *Mathematical Methods of Operations Research*. http://ssrn.com/paper=956223.

[Ye & Tiong, 2000] Ye, S. & Tiong, R. L. K. (2000). NPV-at-risk method in infrastructure project investment evaluation. *Journal of construction engineering and management*, 126(3), 227–233. doi:10.1061/(ASCE)0733-9364(2000)126:3(227).

A Appendix

A.1 Probability measures

A.1.1 Real world probabilities

Real world probabilities are a measure of the actual frequency of an underlying event occurring. The use of real world probabilities allows events to occur as often as they do in the real world. This gives realistic projected returns, allowing risk metrics (see section 2.2) to be calculated. Unfortunately, many methods that use real world projections are not able to produce discount rates that give consistent asset prices.

Numerical Example Imagine an example economy with only two future states, boom and bust, containing two assets with unknown current market price but known future cash flows.

Payoff matrix	Boom state payoff	Bust state payoff
Real world probability	$1/5$	$4/5$
Asset A	3	1
Asset B	2	0.5

By taking expectations based on real world probabilities and using a discount rate of 0.95, the price of each asset can be calculated as

	Implied price
Asset A	1.33
Asset B	0.76

There is nothing wrong with this pricing as all available information is used and, on average, these are fair prices for the assets.

Now consider the case where the current market prices are known to be different to the estimated values implied by the real world probabilities.

Price/Payoff matrix	Current price	Boom state payoff	Bust state payoff
Asset A	1.65	3	1
Asset B	1.00	2	0.5

A new asset is introduced which also has known future cashflows in each state. The current price could be calculated based on the real world probabilities.

	Current price	Boom state payoff	Bust state payoff
Asset C	1.14	2	1

But this price cannot be correct. If 2 of asset A are bought, and 2 of asset B are sold short[12] then a portfolio is constructed with a payoff exactly matching asset C, but with a price of 1.30. An investor selling asset C would be losing 0.16 on every item sold.

[12] Selling short involves borrowing an asset from someone else; selling it; repurchasing the asset once the future state occurs, at the new price; returning it to the person it was borrowed from, allowing profit to be made from a drop in price (or losses from an increase in price)

A stock market trader could construct a portfolio selling short 2 of asset A, buying 2 of asset B and 1 of asset C. The portfolio payoff would be 0 in the future no matter which state occurred, but 0.16 would be received now. The trader would be able to bank this amount every time the portfolio was constructed, and would continue to do so until the prices of assets B or C rose, or asset A fell.

A.1.2 Risk neutral probabilities

The risk neutral probability measure is an artificial construct that simplifies pricing models by adjusting the probabilities of states occurring. Doing this allows the expected returns on any asset, including risky and state dependent assets, to be discounted at the risk free rate.

These market implied risk neutral probabilities show what investors are afraid of happening, not what is likely to happen. Projected prices are not realistic as they are generated under an artificial probability measure, only increasing at the risk free rate. For example, if the market is risk averse then undesirable events will have a disproportionately high probability of occurring. The adjustment of probabilities away from the real world makes some risk measures meaningless.

If a risk neutral modelling environment is already available (i.e. a company may already have a risk neutral derivative pricing framework) then little extra work is required to add risk premium drifts to projections, producing an approximation to the real world process. This adjustment is only exact if normal distributions are assumed for the log of asset prices.

Unfortunately the concept of a risk neutral probability is difficult to understand and therefore hard for management to trust and base decisions upon. Risk neutral measures of distributions displaying skew and kurtosis are tricky to calculate, and especially difficult to communicate. Only one risk measure can be projected at a time, so it is not possible to use real world projections for some assets and risk neutral probabilities for others. The risk free rate and risk neutral law change when currency changes.

Numerical Example By choosing a discount rate of 0.95 and assuming risk neutral probabilities of each state occurring of

	Boom state payoff	Bust state payoff
Risk neutral probability	$7/19$	$12/19$

Then the prices of asset A and asset B match the current market price

Price/Payoff matrix	Current price	Boom state payoff	Bust state payoff
Asset A	1.65	3	1
Asset B	1.00	2	0.5

and the price of asset C matches the arbitrage free price.

	Current price	Boom state payoff	Bust state payoff
Asset C	1.30	2	1

Risk neutral pricing can accommodate (almost) any price for assets A and B, as a suitable choice of discount rate and probabilities can be made to ensure that risk neutral pricing is correct.

A.2 Absence of arbitrage

Arbitrage opportunities offer the ability to make probable future profit, without downside risk and with zero initial investment. Absence of arbitrage implies that a portfolio cannot be created for free that may payoff positively but will certainty never make a loss. It implies that if one asset's cashflows are always greater than another then its price must be higher. It implies that the price of any positive payoffs are also positive.

If two portfolios have the same payoff cashflows in every state of nature then they must have the same price. If this were not true then the cheaper asset could be bought, and the more expensive sold, generating a positive cashflow and forming a portfolio with zero payoffs. This is the law of one price.

Absence of arbitrage means that instant profits cannot be made by repackaging portfolios. If there exists at least one investor that values a package only by contents then arbitrage opportunities will not exist as this investor will repeatedly trade them until they disappear.

A financial model can only construct a meaningful concept of price if it is free of arbitrage. A model that allows arbitrage will give spurious answers if used to optimize investment strategies. Unfortunately it is easy for arbitrage to sneak into models by accident.

Arbitrage Example The following example economy has only two possible future states, boom and bust. In the economy there are two assets, both with a known current market price. The cashflow returned by each asset in each of the economy's states is also known.

Price/Payoff matrix	Current price	Boom state payoff	Bust state payoff
Asset A	1.65	3	1
Asset B	1.20	2	0.5

This pricing structure contains arbitrage. Assume a portfolio can be constructed from any weighting of assets and that there are no frictional costs, no bid offer spread and no short sell limit. If a portfolio is created by buying 8 of asset A, funded by selling short 11 units of asset B then the portfolio will have a zero current value but will guarantee to return at least 2. If prices stay the same then the process can be repeated creating unlimited free money.

In reality, if this situation occurred then the prices would not stay the same. The demand created through buying asset A will drive the price up and the supply created through selling asset B will drive its price down, reducing the profitability of the arbitrage. This repricing would continue until the arbitrage opportunity was removed from the market. Profit opportunities using this sort of trading strategy do not exist frequently or for long in the real world. It quickly becomes too expensive to construct a portfolio that eliminates enough risk to ensure that profit is certain.

A.3 Mean reversion

Mean reversion is a feature of many commodity prices. It can also be used as a way to control the likely upper and lower values of an asset. [Clewlow & Strickland, 2000] describe the regression approach to estimating mean reverting parameters.

For a given set of observations Z_{t_i} of the mean reverting process, viewed at times t_i, the regression equation is $\Delta Z_{t_i} = a_0 + a_1 Z_{t_i} + \varepsilon_{\Delta t}$ where $a_0 = \overline{Z}\left(1 - \exp\left(-\alpha \Delta t\right)\right)$ and $a_1 = \exp\left(-\alpha \Delta t\right) - 1$ allowing the relationship between Δx_t and x_t to be tested in the presence of noise $\varepsilon_{\Delta t}$.

By regressing Δx_t against x_t it is possible to get estimates of α_0 and α_1, which can be rewritten in terms of the parameters $\overline{Z} = -\frac{a_0}{a_1}$ and $\alpha = \frac{-\ln(a_1+1)}{\Delta t}$.

Before any regression analysis the data should be cleaned. This means that the data may need to have growth factors removed or be adjusted for seasonal effects. These adjustments can make choosing the value for $\overline{Z}$ over time more an art than a science.

The mean reversion parameter α cannot be negative. If the regressions produce negative values then it implies that the process is not mean reverting. Low values for α may also imply a non-mean reverting process. If the sample length is low then it is possible for the data to show no mean reversion when some may actually exist.

Half life When describing models it is often useful to refer to a real life measurement of mean reversion, rather than talk in terms of the model parameter α. The half life is defined to be the time taken $t_{1/2}$ for the difference in value between the current level of Z and its long term level $\overline{Z}$, to halve.

The relationship $t_{1/2} = \frac{\ln(2)}{\alpha}$ allows the half life $t_{1/2}$ and the mean reversion parameter α to be used interchangeably.

A.4 Correlation

Control and understanding of covariance is essential for the Monte Carlo SDF framework. Section A.4 discusses how mean reversion affects correlations over time, and shows how to generate correlated random variables. Correlations do not change over time when processes are not mean reverting, but if a process is mean reverting then the correlation with other processes is dependent on the timeperiod being projected.

This section considers two cases of paired processes, one where both processes are mean reverting, and the other where only one of the processes mean reverts.

A.4.1 A mean reverting and a non-mean reverting process

Define Z_t^i as

$$Z_t^i = \overline{Z}^i + \left(Z_0 - \overline{Z}^i\right)\exp\left(-\alpha_i t\right) + \int_0^t \sigma_i \exp\left(-\alpha_i\left(t-r\right)\right) dX_r^i$$

and Z_t^k as $Z_t^k = \mu t + \int_0^t \sigma_k dX_k$ and set $E\left[dX_i dX_k\right] = \rho_{ik} dr$ giving a covariance of

$$\begin{aligned} Cov\left(Z_t^i, Z_t^k\right) &= Cov\left(\int_0^t \sigma_i \exp\left(-\alpha\left(t-r\right)\right) dX_i, \int_0^t \sigma_k dX_k\right) \\ &= E\left[\int_0^t \sigma_i \sigma_k \exp\left(-\alpha\left(t-r\right)\right) dX_i dX_k\right] - 0 \end{aligned}$$

leaving

$$Cov\left(Z_t^i, Z_t^k\right) = \int_0^t \sigma_i \sigma_k \exp\left(-\alpha_i\left(t-r\right)\right)\rho_{ik} dr = \sigma_i \sigma_k \frac{\rho_{ik}}{\alpha_i}\left(1-\exp\left(-\alpha t\right)\right) \tag{45}$$

The correlation over time is

$$\widehat{\rho}_{ik}^t = \frac{Cov\left(Z_t^i, Z_t^k\right)}{\sqrt{V\left[Z_t^i\right] V\left[Z_t^k\right]}} = \rho_{12}\frac{2\left(1-\exp\left(-\alpha t\right)\right)}{\sqrt{t\alpha\left(1-\exp\left(-2\alpha t\right)\right)}}$$

so the long term correlation is $\widehat{\rho}_{ik}^{\infty} = 0$ as $t \to \infty$.

When projecting far in the future, if only one of the processes mean reverts then this destroys any correlation. The variance of the non mean reverting process becomes large in comparison to the mean reverting process, drowning out any correlation.

A.4.2 Market beta for a mean reverting process

The market is usually modelled without mean reversion implying that any asset displaying mean reversion will have a risk premium that declines over time. The reduction in expected return can be demonstrated using the CAPM (see section 6.2.2).

The market beta is defined as $\beta^i = \frac{Cov\left(R_t^i, R_t^p\right)}{V\left[R_t^p\right]}$ where R_t^i is the return on the mean reverting process with instantaneous volatility σ_i and mean reverting parameter α. The variance of the market return is $\sigma_p^2 t = V\left[R_t^p\right]$ and using eq: 45 from the previous section

$$Cov\left(R_t^i, R_t^p\right) = \sigma_i \sigma_p \frac{\rho_{ip}}{\alpha}\left(1 - \exp\left(-\alpha t\right)\right)$$

So

$$\beta_t^i = \frac{\sigma_i}{\sigma_p}\frac{\rho_{ip}}{\alpha t}\left(1 - \exp\left(-\alpha t\right)\right) \tag{46}$$

which decays to 0 over time. Note that if the process is not mean reverting then $\alpha = 0$ and the formula reverts back to the time homogeneous form $\beta_t^i = \frac{\sigma_i}{\sigma_p}\rho_{ip}$. A numerical example of this equation is given in section 6.2.2.

When modelling using the covariance pricing model (section 3.5.1), controlling correlations is essential. The correlation of a mean reverting process with the SDF will be dependent on the timeperiod projected. This implies that the risk premium placed on a mean reverting asset will fade over time.