

The Stochastic Discount Factor for the Exponential-Utility Capital Asset Pricing Model

Mark Johnston
PricewaterhouseCoopers,
GPO Box 2650, Sydney NSW 1171, Australia
and
University of New South Wales,
Sydney NSW 2052, Australia

E-mail: mark.e.johnston@au.pwc.com

May 20, 2004

Abstract

In this paper we introduce the notion of a marginal stochastic discount factor for an asset. This is a function defined on the possible payoffs of the asset which can be used to value the asset and derivatives of it. We then show how one may derive the marginal stochastic discount factor for a common version of the Capital Asset Pricing Model (CAPM), namely the version derived under assumptions of exponential utility and asset payoffs jointly normal with the market portfolio. The advantage of this is that it allows us to obtain equilibrium prices for assets which are derivatives of CAPM assets, even though these derivatives do not satisfy the CAPM.

We then illustrate the practical application of this discount factor by showing how to value the liabilities of a firm with CAPM assets.

Our derivation of the CAPM discount factor builds upon some papers published in the actuarial literature by Shaun Wang and Hans Bühlmann.

1 Introduction

The capital asset pricing model (CAPM), developed in the 1960's, is the standard model employed today to value many assets, such as companies and projects. The CAPM is applied as a *risk-adjusted discount rate* technique. The latter says that an asset with payoff x_j (a random variable) can be valued according to the formula:

$$p_j = \frac{E(x_j)}{E(R_j)}.$$

Here p_j denotes the price of the asset, and the number $E(R_j)$ is the expected or required return on the asset. The CAPM gives a particular formula for the required return in terms of the relative riskiness of the asset and the market. This formula applies only to assets which satisfy the assumptions of the CAPM.

Modern equilibrium asset pricing theory is cast in terms of *stochastic discount factors*. These are random variables that allow assets to be priced according to formulae like:

$$p_j = E(m x_j),$$

where m is the stochastic discount factor (for example see [Coc01, LW01]). The single discount factor m may be used to value all traded assets. The stochastic discount factor approach is very general, which is especially useful for consideration of theoretical valuation questions. To apply the approach to answer practical valuation questions, one can either estimate a discount factor directly from data, as described in Cochrane's book [Coc01], or employ a specialisation of the approach that is reasonable for whatever particular class of assets one is interested in.

In this paper we introduce the notion of a marginal (stochastic) discount factor for an asset. This is a function defined on the possible payoffs of the asset which can be used to value the asset and derivatives of it. We then show how one may derive the marginal stochastic discount factor for a common version of the CAPM, namely the version derived under assumptions of exponential utility and asset payoffs jointly normal with the market portfolio. The advantage of this is that it allows us to obtain equilibrium prices for assets which are derivatives of CAPM assets, even though these derivatives do not satisfy the CAPM.

We then illustrate the practical application of this discount factor by showing how to value the liabilities of a firm with CAPM assets.

Our derivation of the CAPM discount factor builds upon some papers published in the actuarial literature by Shaun Wang and Hans Bühlmann. In his paper of 1980 [Büh80a], Hans Bühlmann developed a model for de-

termining “optimal risk exchanges” in a closed market. Wang [Wan03] has recently shown that the “Wang Transform,” originally developed in a probability distortion operator approach to asset pricing [Wan02], arises also in an expected-utility equilibrium context, in a specialisation of Bühlmann’s model. What we think is noteworthy about the Bühlmann paper is that it provides an explicit formula for what in modern terminology is the stochastic discount factor, as a function of *total* consumption.

Here we will set out the results of Bühlmann in the more familiar framework of expected-utility asset pricing, and then will go on to show how Wang’s transform arises as a way of pricing certain assets within a specialisation of the Bühlmann economy. For these assets, which we might call the Wang assets, we can exhibit an explicit formula for the marginal stochastic discount factor. We will then show that a subset of the Wang assets, those which have normally distributed payoffs, can be priced using the capital asset pricing model. This set of assets is closed under linear combinations, and so forms a linear subspace of the space of all assets in the Wang-Bühlmann economy. We exhibit an explicit form for the marginal discount factor for CAPM assets — which turns out to be exponential in the payoff of the asset.

2 Background

In this section we will introduce the stochastic discount factor approach to asset pricing, and describe the connection between this approach and the risk-neutral valuation approach. We will discuss returns and Sharpe ratios, and show how Sharpe ratios of assets may be related to the Sharpe ratio of the stochastic discount factor. We will set out some formulae for expected returns and Sharpe ratios of linear combinations of assets, which will be useful in what follows.

2.1 Securities market equilibrium

In modern texts asset pricing texts such as LeRoy and Werner [LW01], Cochrane [Coc01] and Panjer *et. al.* [P⁺98], it is shown that at a complete-market equilibrium in a single-period expected-utility endowment economy, an asset with payoff x may be priced via the formula:

$$p = E(mx),$$

where $m = m_i$ for all i , and

$$m_i = \frac{u'_{i,1}(c_{i,1}^*)}{u'_{i,0}(c_{i,0}^*)} \tag{1}$$

is agent i 's *marginal rate of substitution* between consumption at the start of the period and consumption at the end of the period. At the complete markets equilibrium, the agents' marginal rates of substitution agree, and we refer to m as the *stochastic discount factor*. Here we have supposed that agent i has a utility function of the form (“separable”): $u_i(c_0, c_1) = u_{i,0}(c_0) + u_{i,1}(c_1)$, where c_t represents the agent's consumption at time t . The equilibrium forms through agents buying and selling securities with the aim of maximising their expected utility. We denote the resulting optimal consumption allocations by $\{c_{i,0}^*, c_{i,1}^*\}$. The optimal end-of-period consumption $c_{i,1}^*$ is a random variable — it varies over “states of the world”. We can think of states of the world as being elements ω of a state space Ω , and random variables as functions from Ω to the real numbers \mathbf{R} . Since $c_{i,1}^*$ is a random variable, the stochastic discount factor m is a random variable.

At equilibrium, the securities market clears — for every buyer there is a seller. Writing $\theta_{i,j}^*$ for the optimal amount of security j held by agent i , the market clearing conditions are:

$$\sum_{i=1}^I \theta_{i,j}^* = 0,$$

for each security j . We will write the total end-of-period consumption as $c(\omega) = \sum_{i=1}^I c_{i,1}^*(\omega)$. In a single-period model, end-of-period consumption equals end-of-period wealth [Coc01, P+98]. We may thus consider c to be total wealth.

2.2 No-arbitrage

A set of payoffs and their corresponding prices are said to be *arbitrage-free* if each positive payoff has a positive price. By *positive payoff* we mean a payoff that is not negative in any state, and is positive in at least one state. By *strictly positive payoff* we mean a payoff which is positive in all states.

This property of prices can be translated into a property of the discount factor: A set of payoffs and prices is arbitrage-free if and only if there exists a strictly positive discount factor (see e.g. [Coc01]).

2.3 Relationship with risk-neutral pricing

Here we describe briefly the connection between pricing using stochastic discount factors and pricing using *risk-neutral probabilities*.

If we make the assumptions of no-arbitrage, and complete markets, the stochastic discount factor approach becomes the *risk-neutral pricing* approach. This is easiest to think about in the finite-state-space context, so

assume Ω is a finite set, and let $\pi(\omega)$ denote the probability of state ω . Then the price of an asset that pays $x(\omega)$ in state ω is:

$$p = E(mx) = \sum_{\omega \in \Omega} m(\omega) x(\omega) \pi(\omega) = E(m) \sum_{\omega \in \Omega} x(\omega) \pi^Q(\omega),$$

where we have defined the numbers $\pi^Q(\omega)$ as

$$\pi^Q(\omega) = \pi(\omega) \frac{m(\omega)}{E(m)}. \quad (2)$$

The numbers $\pi^Q(\omega)$ add up to one, by construction, and since we have assumed no-arbitrage, the discount factor m is strictly positive (see e.g. [Coc01]), so the numbers $\pi^Q(\omega)$ are non-negative. Thus we may consider them a set of probabilities, which we will call the *risk-neutral probabilities*, and write:

$$p = \frac{E^Q(x)}{R_f}, \quad (3)$$

where we have defined $R_f = 1/E(m)$, and E^Q represents the expectation with respect to the risk-neutral probabilities, rather than the original, “physical”, probabilities (which give the likelihoods of states of the world). The labelling of $1/E(m)$ as R_f follows from the fact that if a risk-free asset is traded, meaning an asset which pays an amount 1 in every state, then its price is $E(m)$, so $1/E(m)$ can be interpreted as the risk-free gross return. Equation (3) embodies the risk-neutral pricing approach: “To value an asset, calculate its expected payoff with respect to the risk-neutral probabilities, and discount at the risk-free rate”.

Equation (2) shows that the stochastic discount factor describes the transformation between physical and risk-neutral probabilities — the stochastic discount factor, scaled by its mean, may be thought of as the derivative of the risk-neutral probabilities with respect to the physical probabilities, or a “change of measure”.

2.4 Returns, and the market price of risk

By the *gross return*, or just *return* on an asset x with non-zero price, we mean the random variable $R_x = x/p(x)$, payoff divided by price. The return is just a scaled version of the asset. It has price 1. Any asset with price 1 can be considered to be a return. The expected return is the number $E(R_x) = E(x)/p(x)$. This expectation is of course taken with respect to the

real-world probability measure. We can easily show that the expected return of any asset in the risk-neutral measure is the risk-free return:

$$E^Q(R_x) = E^Q\left(\frac{x}{p(x)}\right) = \frac{E^Q(x)}{p(x)} = \frac{E^Q(x)}{E^Q(x)/R_f} = R_f.$$

In the real-world probability measure one finds, through application of the pricing formula $p = E(mx)$ and the definition of covariance, that

$$E(R_x) = R_f(1 - \text{cov}(m, R_x)), \quad (4)$$

so that expected returns are greater/lesser than the risk-free return according to whether the covariance of the asset with the stochastic discount factor is negative/positive, or equivalently (when the SDF has been derived in an equilibrium model) whether the covariance of the asset with consumption is positive/negative.

For any asset x with non-zero price we define the *Sharpe Ratio* as the ratio of expected excess return to standard deviation of return:

$$\lambda_x = \frac{E(R_x) - R_f}{\sigma_{R_x}}. \quad (5)$$

This ratio is also known as the *market price of risk of asset x* , as it shows how much excess return will be demanded per unit of risk (measured as return standard deviation). Applying formula (4) yields:

$$\lambda_x = \frac{-R_f \text{cov}(m, R_x)}{\sigma_{R_x}}.$$

Multiplying top and bottom by $p(x)$ gives:

$$\lambda_x = \frac{-R_f \text{cov}(m, x)}{\sigma_x},$$

where σ_x is the standard deviation of x . Accordingly, we have:

$$\lambda_x = -R_f \sigma_m \rho_{mx},$$

where σ_m is the standard deviation of the stochastic discount factor, and ρ_{mx} is the linear correlation coefficient of m and x . We can consider m as an asset, in which case

$$\lambda_m = \frac{-R_f \text{cov}(m, m)}{\sigma_m} = -R_f \sigma_m,$$

so for any asset x we have:

$$\lambda_x = \rho_{mx} \lambda_m. \quad (6)$$

In words: the Sharpe Ratio for an asset x is the product of the Sharpe Ratio of the stochastic discount factor and the linear correlation coefficient of the asset with the discount factor. The reader may note that if m is interpreted as “the market”, equation (6) is the capital asset pricing model. We will discuss this further later. In general m does not admit such an interpretation, but this CAPM-like formula holds.

We’ll now show how expected returns and Sharpe Ratios behave when we form portfolios of assets. Suppose an asset x is the sum of two assets y and z : $x = y + z$. Then:

$$p(x) R_x = p(y) R_y + p(z) R_z,$$

so:

$$R_x = \frac{p(y)}{p(x)} R_y + \frac{p(z)}{p(x)} R_z,$$

and consequently the expected return of x is given by:

$$E(R_x) = \frac{p(y)}{p(x)} E(R_y) + \frac{p(z)}{p(x)} E(R_z).$$

This is a “weighted average cost of capital” formula. The weights are the relative market values of y and z .

Using the definition of the Sharpe Ratio in terms of the mean and standard deviation of the payoff of x , we see that:

$$\begin{aligned} \sigma_x \lambda_x &= -R_f \operatorname{cov}(m, x) \\ &= -R_f \operatorname{cov}(m, y + z) \\ &= -R_f ((\operatorname{cov}(m, y) + \operatorname{cov}(m, z))) \\ &= \sigma_y \lambda_y + \sigma_z \lambda_z, \end{aligned}$$

and hence we obtain a “weighted average price of risk” formula:

$$\lambda_x = \frac{\sigma_y}{\sigma_x} \lambda_y + \frac{\sigma_z}{\sigma_x} \lambda_z,$$

where the weights are the relative standard deviations of the asset payoffs.

3 Marginal discount factors

A stochastic discount factor is a function from Ω to \mathbf{R} . The state space Ω is in general a high-dimensional space. Often we will want to use a stochastic discount factor to price a particular asset, or derivatives of a particular asset, and we'd like to graph the stochastic discount factor as a function of the payoff of the asset. When we do this, we are actually graphing what we might call the *marginal stochastic discount factor* for that asset. To make this concrete, suppose x is an asset, m is a stochastic discount factor, and for each $s \in \mathbf{R}$ let $x^{-1}(s)$ denote the set of states $\{\omega \in \Omega : x(\omega) = s\}$ (the *level set* of x corresponding to the value s). Consider a *derivative* of x , meaning an asset with payoff of the form $g(x)$ for some function g . We call it a derivative of x because its payoff is determined by the payoff of x . Writing f for the probability density function corresponding to the real-world probabilities, the value of this derivative is given by:

$$\begin{aligned}
 E(m g(x)) &= \int_{\omega \in \Omega} m(\omega) g(x(\omega)) f(\omega) d\omega \\
 &= \int_{s \in \mathbf{R}} \int_{\omega \in x^{-1}(s)} m(\omega) g(s) f(\omega) d\omega ds \\
 &= \int_{s \in \mathbf{R}} g(s) f_x(s) \left(\int_{\omega \in x^{-1}(s)} m(\omega) \frac{f(\omega)}{f_x(s)} d\omega \right) ds \\
 &= \int_{s \in \mathbf{R}} m_x(s) g(s) f_x(s) ds \\
 &= E(m_x g(x)),
 \end{aligned}$$

where $m_x(s)$ is the quantity in brackets, and f_x is the marginal density of x . The quantity $f(\omega)/f_x(s)$ is the conditional density with respect to x . Thus m_x is the conditional expectation of the SDF with respect to x .

As m_x is defined in the same way as the marginal probability density of x , that is by integrating over the level sets of x , we refer to m_x as the *marginal discount factor* of x . Just as the marginal density is all one needs to calculate expectations of functions of x , so the marginal density and marginal discount factor are all one needs to compute the values of functions (“derivatives”) of x . When we draw a graph of a discount factor it's usually a marginal discount factor that we are drawing.

4 Derivation of Bühlmann's model in the SDF framework

Under the standard assumption that marginal utility is strictly positive, the end-of-period marginal utility function will have a local inverse at the optimal consumption point, so we may invert the relationship (1) to obtain:

$$c_{i,1}^*(\omega) = (u'_{i,1})^{-1} (u'_{i,0}(c_0^*) m_i(\omega)).$$

We can thus express total consumption in terms of the discount factors as follows:

$$c(\omega) = \sum_{i=1}^I c_{i,1}^*(\omega) = \sum_{i=1}^I (u'_{i,1})^{-1} (u'_{i,0}(c_0^*) m_i(\omega)). \quad (7)$$

To proceed further we need to assume that we can replace the m_i with m in this formula. As noted above, this will be possible in a complete markets setting. To deal with the incomplete markets case we need to understand what happens when we take projections on each side of the above equation. The left hand side would become the traded part of consumption/wealth (call it the “market portfolio”). We will return to this matter in a subsequent paper.

Replacing each m_i with m in (7), we have:

$$c(\omega) = \sum_{i=1}^I (u'_{i,1})^{-1} (u'_{i,0}(c_0^*) m(\omega)). \quad (8)$$

In Bühlmann's second paper on this topic [Büh80b], he notes that in the general case one can use the implicit function theorem to argue that equation (8) can be solved for m in terms of c . In his first paper, he obtains an explicit solution by assuming that agents have exponential utility functions:

$$u_{i,1}(c) = \frac{1}{\alpha_i} (1 - e^{-\alpha_i c}),$$

where α_i is a parameter known as the *absolute risk-aversion*, in which case:

$$u'_{i,1}(c) = e^{-\alpha_i c},$$

and so:

$$(u'_{i,1})^{-1}(s) = \frac{-1}{\alpha_i} \log s.$$

Under this exponential utility assumption the terms of the sum in (8) split and we get:

$$c(\omega) = \sum_{i=1}^I -\frac{1}{\alpha_i} \log(u'_{i,0}(c_0^*)) - \left(\sum_{i=1}^I \frac{1}{\alpha_i} \right) \log(m(\omega)).$$

We then conclude that m is an exponential function of aggregate consumption

$$m(\omega) = A e^{-\alpha c(\omega)},$$

where α is the harmonic sum of the α_i :

$$\frac{1}{\alpha} = \sum_{i=1}^I \frac{1}{\alpha_i},$$

and A is a positive constant. A formula for the constant A can be written down in terms of the risk-aversion coefficients and the time-zero marginal utilities, or if a risk-free asset is traded we can calibrate A by requiring that the expected value of the discount factor be the inverse of the risk-free return.

5 Stochastic discount factor for Wang assets

In the previous section we showed that in the Bühlmann economy we could obtain an explicit form for the stochastic discount factor, as a function of aggregate consumption. In this section we will place a further restriction on the economy, namely that aggregate consumption is normally distributed, and will show that this allows us to determine marginal discount factors for certain assets (those which are *normal-copula* with consumption).

Accordingly, following Wang [Wan03], assume c is normally distributed, and let z be the standard normal variable obtained from c by subtracting its mean and dividing by its standard deviation: $z = (c - \mu_c)/\sigma_c$. Then the discount factor can be re-written as $m = B e^{-\lambda z}$, where $\lambda = \alpha \sigma_c$. The stochastic discount factor thus has a log-normal distribution in this economy.

Assume an asset x is normal-copula with total consumption, meaning here that x may be written as $x = h(v)$, where h is an increasing function, and v and z have a standard bi-variate normal distribution with correlation coefficient ρ . In this case we can decompose z in a way which allows us to directly compute what the marginal discount factor for x . Let $y = z - \rho v$. Then y and v are also jointly normal, and as $\text{cov}(y, v) = \text{cov}(z, v) - \rho \text{cov}(v, v) = 0$, y is independent of v (for jointly-normal variables, zero covariance implies independence).

Consider a derivative of x , $g(x)$. The value of this asset is given by:

$$E(m g(x)) = E(B e^{-\lambda(y+\rho v)} g(h(v))) = E(B e^{-\lambda y} e^{-\lambda \rho v} g(h(v)))$$

Since y and v are independent, we may write this as:

$$E(m g(x)) = E(B e^{-\lambda y}) E(e^{-\lambda \rho v} g(h(v))) = E(m_x g(x)),$$

say, where

$$m_x = C e^{-\lambda \rho h^{-1}(x)},$$

and C is a constant. Thus we have derived a stochastic discount factor, m_x , which prices any derivative of x via the one-dimensional integral:

$$p(g(x)) = E(m_x g(x)) = \int_{-\infty}^{\infty} m_x(s) g(s) f_x(s) ds,$$

where f_x is the marginal probability density of x . The function m_x is the marginal discount factor for x , as per our earlier definition.

The fact that the variable v is jointly-normal with consumption means that all information about the dependence between v and consumption is captured by a single number, the correlation coefficient. This, together with the assumption of exponential utility, has allowed us to separate the expression $E(m g(x))$ into a product of two expectations, which means two integrals, one over x and one over all the other dimensions of the state space, unconditional on x . This is what gives us an explicit form for the marginal discount factor in this case.

When a risk-free asset is traded, the constant C above may be determined by requiring that the expected value of m_x be $1/R_f$, in which case we find that:

$$m_x(s) = \frac{1}{R_f} e^{-\frac{\lambda_x^2}{2}} e^{-\lambda_x h^{-1}(s)}, \quad (9)$$

where $\lambda_x = \lambda \rho$. Wang refers to the parameter λ_x as the *market price of risk for the asset x* , and we will discuss later the reasons for this. It is really the market price of risk of the underlying asset v .

5.1 Connection with the Wang Transform

In the remainder of the paper [Wan03], Wang goes on to show that the prices of Wang assets may also be calculated as:

$$p(x) = \frac{1}{R_f} E^Q(x),$$

where E^Q represents the expectation with respect to a transformed probability distribution F_x^Q , the *Wang Transform* of F_x , defined by

$$F_x^Q(s) = \Phi(\Phi^{-1}(F_x(s)) + \lambda_x), \quad (10)$$

where F_x is the cumulative distribution function (CDF) of x and Φ is the CDF of a standard normal variable.

Note that by differentiating Equation (10) we can establish a formula for the probability density function of the transformed variable:

$$f_x^Q(s) = (F_x^Q)'(s) = e^{\frac{-\lambda_x^2}{2}} f_x(s) e^{-\lambda_x \Phi^{-1}(F_x(s))},$$

where $f_x(s)$ is the original probability density function: $f_x(s) = F_x'(s)$. We may then apply Equation (2) to recover the formula (9) for the stochastic discount factor (where $h^{-1}(s) = \Phi^{-1}(F_x(s))$).

6 Stochastic discount factor for CAPM

We have shown that in a single-period endowment economy, where agents are expected-utility maximisers and have separable utility functions, with the end-of-period utility being exponential, and where the total end-of-period endowment/consumption is normally distributed, we can derive an explicit expression for the stochastic discount factor in terms of total consumption, and that for certain assets within this economy, namely those which are normal-copula with consumption, we can derive an explicit expression for the marginal stochastic discount factor.

The set of assets which are jointly normal with consumption are of particular interest, as we shall see. For brevity we'll call them *normal assets*, or *CAPM assets* (see below) but remember that this means more than just that they have a normal distribution, it means they are jointly normal with consumption. This is a stronger condition — two variables can each have normal marginal distributions but not be jointly normal; their dependency structure could be characterised by some other copula (see e.g. [Ven03]).

If x and y are jointly normal with consumption, then so is $ax + by$ for any real numbers a and b , so the space of normal assets is closed under linear combinations, and so forms a vector subspace of the space of all assets. This means that a portfolio of normal assets is a normal asset. This property does not hold for the Wang assets, for example.

Actually, since the marginal stochastic discount factor for an asset can be used to price all derivatives of that asset, and since any asset which is normal-copula with consumption is a derivative of an asset which is jointly

normal with consumption ($x = h(v)$ in the above), we need only consider assets which are jointly normal with consumption to derive the benefits of the Wang transform in this expected-utility setting. However, it may be more convenient to work with the marginal distribution and marginal discount factor of the asset under investigation, than to work in terms of those of the underlying normal asset.

If an asset x has payoff which is normal with mean μ_x and standard deviation σ_x , then it may be related to a unit normal variable v via the increasing transformation $x = h(v) = \mu_x + \sigma_x v$. The inverse transformation is $h^{-1}(x) = (x - \mu_x)/\sigma_x$, so we see from (9) that the marginal stochastic discount factor for x is:

$$m_x(s) = \frac{1}{R_f} e^{-\frac{\lambda_x^2}{2}} e^{-\lambda_x (s - \mu_x)/\sigma_x}, \quad (11)$$

so the marginal stochastic discount factor for a normal asset is exponential in the payoff of the asset.

6.1 Relationship with traditional CAPM

The reader may recall that assumptions of exponential utility and asset payoffs jointly normal with consumption give one way of deriving the Capital Asset Pricing Model (CAPM) (see e.g. [Coc01]). Normally one arrives at the CAPM by assuming that all assets have payoffs jointly normal with consumption, and proceeding by pricing assets relative to the “market portfolio” consisting of all assets in the economy. Here we have taken a slightly different approach, and allowed in the securities market assets which have non-normal payoffs. What we have shown is that we can obtain a stochastic discount factor that prices the normal assets, and derivatives of them. We can thus price a broader set of assets, using the normal assets as a base, and employing the two tools at our disposal, namely the ability to price derivatives of normal assets, and the ability to price portfolios of assets, which follows from the linearity of the pricing operator.

We have yet to show that our marginal discount factors applied to the normal assets produce the CAPM. We’ll see this below.

6.2 Properties of the CAPM discount factor

Suppose x is a derivative of a normal asset v : $x = h(v)$. Without loss of generality we can assume v is unit normal (a derivative of a normal asset is also a derivative of a unit normal asset). Then we may value x using the

marginal discount factor of v as follows:

$$\begin{aligned}
p(x) &= E(m_v h(v)) \\
&= \int_{-\infty}^{\infty} m_v(s) h(s) f_v(s) ds \\
&= \int_{-\infty}^{\infty} \frac{1}{R_f} h(s) e^{-\frac{\lambda_v^2}{2}} e^{-\lambda_v s} e^{\frac{-s^2}{2}} \frac{1}{\sqrt{2\pi}} ds \\
&= \int_{-\infty}^{\infty} \frac{1}{R_f} h(s) \frac{e^{-\frac{(s+\lambda_v)^2}{2}}}{\sqrt{2\pi}} ds \\
&= \int_{-\infty}^{\infty} \frac{1}{R_f} h(s - \lambda_v) \frac{e^{-\frac{s^2}{2}}}{\sqrt{2\pi}} ds \\
&= \frac{1}{R_f} E(h(v - \lambda_v))
\end{aligned}$$

This is a convenient formula for pricing normal assets and their derivatives. We'll consider a couple of examples below. Note that the λ used is that of v , i.e. $\lambda_v = \rho_{vc} \lambda$.

Value of a normal asset Let x be a normal asset with mean μ and standard deviation σ . Then x is a derivative of a unit normal asset v via $x = h(v) = \mu + \sigma v$, so the price of x is:

$$p(x) = \frac{1}{R_f} E(\mu + \sigma(v - \lambda_v)) = \frac{1}{R_f} ((\mu - \lambda_v \sigma) + \sigma E(v)) = \frac{1}{R_f} (\mu - \lambda_v \sigma).$$

As x is a linear function of v , it has the same correlation with total consumption, so $\lambda_v = \rho_{vc} \lambda = \rho_{xc} \lambda = \lambda_x$.

Value of an asset whose log is normal Let $x = h(v) = e^{\mu + \sigma v}$, where v is a unit normal asset. Then the price of x is:

$$p(x) = \frac{1}{R_f} E(e^{\mu + \sigma(v - \lambda_v)}) = \frac{1}{R_f} e^{\mu - \lambda_v \sigma} E(e^{\sigma v}) = \frac{1}{R_f} e^{\mu - \lambda_v \sigma + \sigma^2/2}.$$

Value of a call option on a normal asset Let $x = \max(0, y - K)$, where y is a normal asset: $y = \mu + \sigma v$. Then $x = \max(0, \mu + \sigma v - K)$, so the price of x is:

$$p(x) = \frac{1}{R_f} E(\max(0, \mu - K + \sigma(v - \lambda_v))) = \frac{\sigma}{R_f} E(\max(0, a + v)),$$

where $a = (\mu - \lambda_v \sigma - K)/\sigma = (R_f p(y) - K)/\sigma$, so

$$p(x) = \frac{\sigma}{R_f} \int_{-a}^{\infty} (a + s) \phi(s) ds = \frac{\sigma}{R_f} (a \Phi(a) + \phi(a)), \quad (12)$$

where ϕ and Φ are respectively the probability density function and cumulative distribution function of a standard normal variable.

6.3 Risk-neutral probabilities for CAPM assets

Applying the definition of risk-neutral probabilities (2), and the form of the marginal SDF for CAPM assets (11) we see that the risk-neutral probability distribution for a normal asset x is just the physical distribution shifted to the left by a distance $\lambda_x \sigma_x$:

$$f_x^Q(s) = f_x(s + \lambda_x \sigma_x).$$

Wang [Wan02] derives this result using the Wang transform.

6.4 Showing that normal assets satisfy the CAPM

We have been referring to the normal assets as CAPM assets, but we haven't yet shown that they satisfy the CAPM. We'll do that here.

Suppose x is a normal asset with parameters μ_x , σ_x , λ_x . Then the return on x is also a normal asset, with mean $E(R_x) = \mu_x/p(x)$ and standard deviation $\sigma_{R_x} = \sigma_x/p(x)$. Recall that for a normal asset,

$$p(x) = \frac{1}{R_f} (\mu_x - \lambda_x \sigma_x).$$

Solving this for λ_x we get:

$$\lambda_x = \frac{\mu_x - R_f p(x)}{\sigma_x} = \frac{E(R_x) - R_f}{\sigma_{R_x}}.$$

Comparing this to equation (5), we see that for normal assets the SDF parameter λ_x is the Sharpe Ratio of the asset x (which is why we denoted it by λ_x and have been referring to it as the market price of risk of asset x).

Recall that in the Wang-Bühlmann economy total consumption, c , is normal, and has $\lambda_c = \lambda$, so applying the above formula to total consumption yields:

$$\lambda = \frac{E(R_c) - R_f}{\sigma_{R_c}}.$$

We may thus interpret λ as the Sharpe Ratio of consumption, and refer to it as the *market price of risk*.

Recalling that the market price of risk of a normal asset x is related to the overall market price of risk via $\lambda_x = \rho_{xc} \lambda$, where ρ_{xc} is the correlation coefficient of x and c , we see that:

$$\frac{E(R_x) - R_f}{\sigma_{R_x}} = \rho_{xc} \frac{E(R_c) - R_f}{\sigma_{R_c}},$$

and we may re-write this as:

$$E(R_x) = R_f + \beta_x (E(R_c) - R_f),$$

where

$$\beta_x = \rho_{xc} \frac{\sigma_{R_x}}{\sigma_{R_c}}.$$

This is the capital asset pricing model (see e.g. [P⁺98]). As pointed out earlier, in a single-period model total end-of-period consumption equals total end-of-period wealth, so the return on consumption R_c may be thought of as the return on wealth, or the return on “the wealth portfolio”, which is how the CAPM is usually presented. If markets are complete (as we have assumed in deriving the Wang-Bühlmann model), then the wealth portfolio can be represented by a particular combination of traded securities, which we might call “the market portfolio”. From the relationship $\lambda_x = \rho_{xc} \lambda$, and the fact the correlation coefficients are no larger than one in absolute value, we conclude that normal assets can have a Sharpe ratio no greater than that of the market portfolio — so the market portfolio is efficient among normal assets in the sense that it offers the greatest ratio of excess return to risk.

7 Example — a limited-liability firm

Here we will introduce a simple model of a limited-liability firm, following Sherris [She04] and Johnston [Joh03], and illustrate how the discount factors developed above can be applied to value the components of the firm’s balance sheet. The discount factors developed above are useful in this context, as balance sheet components are sums of other balance sheet components, so if we assume certain components (e.g. the assets and the claims) are normal, others (e.g. the surplus) will be also. We first introduce the model without any distributional assumptions, and in the subsequent section we consider the case of a normal surplus.

Suppose the firm exists from time 0 to time T . The assets of the firm pay A_T at time T (a random variable), and have value A_0 at time 0. Suppose the

firm has creditors, being providers of finance or trade creditors for example, who have been promised, and will claim, an amount C_T at time T . Since the firm operates under limited liability, the amount that the creditors will actually receive at time T , which we'll call the *liability payoff*, denoted by L_T , is the lesser of their claim and the available assets:

$$L_T = \min(C_T, A_T).$$

Note that if A_T can take negative values, then creditors may receive a negative payoff at time T . In reality creditors will not pay money into the company in the event of default, but these negative payoffs correspond with situations where the company not only has no assets but has unfunded obligations, for example unfunded employee entitlements, or unfunded environmental clean-up obligations, and those owed will have to meet these obligations from their own pockets, for example the local government may have to undertake environmental remediation work that the company should have paid for.

Let $D_T = C_T - L_T$ be the *deficit* between what creditors were promised and what they will be paid. Then:

$$D_T = \max(C_T - A_T, 0), \tag{13}$$

and $L_T = C_T - D_T$. As for the assets, we'll let the subscript zero denote value at time 0 — so for each balance sheet component X we have $X_0 = E^Q(X_T)/R_f$, or equivalently $X_0 = E(m X_T)$. By the linearity of the pricing operator we have: $L_0 = C_0 - D_0$ — “the value of the liabilities equals the value of claims less the value of the deficit”. Sherris [She04] describes D_0 as the value of the *insolvency exchange option*, as it is the value that accrues to shareholders from the fact that at time T they have the option to exchange the assets for the claims.

Any money remaining after creditors have been paid will accrue to the owners of the firm — the equity providers. The cash flow to equity is thus:

$$E_T = A_T - L_T = A_T - \min(C_T, A_T) = \max(A_T - C_T, 0).$$

Also,

$$E_T = A_T - (C_T - D_T) = (A_T - C_T) + D_T = S_T + D_T,$$

where we denote the difference between the asset payoff and the claims by S_T , the *surplus*. Accordingly, $E_0 = S_0 + D_0$: “equity value is the value of assets less the value of creditors’ claims, plus the value of the insolvency exchange option”.

The balance sheet of the company, at market values, thus looks like that shown in Table 1.

Assets	A_0	$L_0 = C_0 - D_0$	Liabilities
		$E_0 = (A_0 - C_0) + D_0$	Shareholders' Equity

Table 1: Opening balance sheet of a simple limited-liability firm (at market values)

7.1 Financial Strength

One might think that liabilities are just assets with a negative sign in front of them, and this is true in a sense, but what distinguishes liabilities is that as cash flows out of the firm they are subject to the risk that the firm defaults. From the point of view of the creditors, their claims are subject to credit risk, being the risk that this firm defaults and they end up being paid less than they were promised. While it also represents a cash flow out of the firm, and appears on the liability side of the balance sheet, the shareholders' claim is different in that nothing was promised, so there's no notion of a risk of being paid less than promised.

Creditors will obviously be interested in the chance of them being paid less than they were promised (the *probability of default* on their claim), and in how much they will lose, relative to what was promised, if default occurs (the *loss given default*). These topics are the subject matter of credit analysis, practiced by creditors (e.g. lending banks), and their agents (e.g. credit rating agencies). They should also consider how much they care about loss due to default, as captured by their stochastic discount factor. These three factors come together to determine the value of the liabilities. Traditionally, credit risk analysis has focussed on the first two factors, and quality of claims has been characterised by credit ratings (e.g. the "AA" and "BBB" ratings given to debt issues by ratings agencies such as Standard and Poors). More recently, as credit risk has become a traded risk, market value-based measures of credit risk (such as credit spreads) have come to supplement the traditional measures.

How should we characterise the quality of claims in the setup we have outlined above? A claim will be high quality if the value of the corresponding liability is close to the value of the claim, or if the value of the deficit is small, relative to the value of the claim. Accordingly, we introduce the *credit discount ratio*:

$$\zeta = \frac{D_0}{C_0}$$

as a value-based measure of the financial strength of a company, or equivalently of the quality of the company's aggregate liabilities. For example a

ratio of 0.05 means that creditors will apply a discount of 5% to the value of the amount they were promised by the firm. Small ζ corresponds with high financial strength — a “strong balance sheet” (strong in a value-based sense — conceivably a firm could have low ζ by having a relatively high expected loss, but where the loss tended to occur in high-consumption states).

Can we relate ζ to a traditional measure of credit-worthiness, such as credit spread? To indicate how this could be done, consider a simple case where the liabilities consist of a single zero-coupon bond, meaning a promise to pay a fixed amount C_T at time T (so C_T is the “face value” of the bond). For a zero-coupon bond of maturity T , face value C_T and market value L_0 , the credit spread is the number c satisfying:

$$L_0 = e^{-(r_f+c)T} C_T,$$

where r_f is the continuously-compounding risk-free rate for maturity T ($R_f = e^{r_f T}$). Solving for the spread gives:

$$c = \frac{1}{T} \ln \left(\frac{e^{-r_f T} C_T}{L_0} \right).$$

Now if the bond was not subject to the default risk of the firm, its value would be: $e^{-r_f T} C_T$, so in our notation, the value of the claim is:

$$C_0 = e^{-r_f T} C_T.$$

Accordingly,

$$c = \frac{1}{T} \ln \left(\frac{C_0}{L_0} \right),$$

so the relationship between credit spread and credit discount ratio is:

$$c = \frac{-1}{T} \ln(1 - \zeta).$$

Conversely,

$$\zeta = 1 - e^{-cT}.$$

The relationship between credit spread and credit discount given here is actually for all the claims on the firm, rather than just a debt claim. The reason we make a distinction is that the assets of the firm might have the possibility of negative payoff (e.g. if they are normally distributed), while debt claims are typically bounded below at zero — the worst outcome for the debt providers is that they get none of their money back. Accordingly, a debt- and equity-financed firm with assets that can have negative payoffs

will have a residual claim, to the negative part of the asset payoff, which is implicitly borne by employees or other stakeholders. We'll address the pricing of the debt claim alone in a subsequent paper.

If the liabilities of a firm did indeed consist of a single zero-coupon bond then we could just use the traditional measures such as credit spread to characterise the financial strength of the firm. However, the above framework sets out a measure of financial strength which applies to liabilities other than debt (for example insurance claims, environmental rehabilitation obligations), and relates it to the measures applied to debt. The credit spread is expressed as a per-annum measure — the extra return that investors demand per annum for taking on the firm's credit risk — while the credit discount ratio is expressed as a discount to value. Whether a per-annum measure or an absolute measure is better depends upon the circumstances, and upon whether one believes that a risk compounds over time or not. The credit spread needs to be expressed together with a maturity in order for the value of the liabilities to be calculated from those of the claims.

8 Example — a firm with normal surplus

In this section we will illustrate the use of the stochastic discount factor developed above for normal assets, by applying it to compute the default discount for a firm with normal surplus. One example of a firm with a normal surplus is a firm with normal assets and known claims. Another is a firm with normal assets and normal claims. Accordingly, let us assume we are in the Wang-Bühlmann economy and that the surplus of the firm is normal, in the sense introduced above. The surplus is then described by three parameters: the mean and standard deviation of its payoff, μ_S and σ_S , and a market price of risk (Sharpe Ratio), λ_S . If we know the overall market price of risk, λ , and the correlation between S_T and consumption, ρ_{Sc} , then we can compute the market price of risk for the assets as $\lambda_S = \rho_{Sc} \lambda$. The value of the surplus is:

$$S_0 = (\mu_S - \lambda_S \sigma_S) / R_f.$$

As captured by Equation (13), the default deficit D_T is a strike-zero call option on the negative of the surplus, $-S_T = C_T - A_T$. In the present case the negative of the surplus is a normal asset with mean $-\mu_S$ and standard deviation σ_S , and Sharpe Ratio $-\lambda_S$. Therefore the value of the default deficit can be computed from the formula (12), to obtain:

$$D_0 = \frac{\sigma_S}{R_f} (a \Phi(a) + \phi(a)),$$

where

$$a = \frac{-\mu_S + \lambda_S \sigma_S - 0}{\sigma_S} = \frac{-S_0 R_f}{\sigma_S}.$$

The credit discount ratio is then:

$$\zeta = \frac{D_0}{C_0} = \frac{\sigma_S}{R_f C_0} \phi\left(\frac{-S_0 R_f}{\sigma_S}\right) - \frac{S_0}{C_0} \Phi\left(\frac{-S_0 R_f}{\sigma_S}\right).$$

To investigate the behaviour of ζ , make the substitutions:

$$\begin{aligned}\mu' &= \frac{S_0}{C_0}, \\ \sigma' &= \frac{\sigma_S}{R_f C_0}, \\ \delta &= \frac{\mu'}{\sigma'} = \frac{S_0 R_f}{\sigma_S}.\end{aligned}$$

Here we are expressing the surplus value as a multiple of the value of the claims, and also dividing the standard deviation of the surplus payoff by R_f to make it a present value and thus comparable to the other quantities (S_0 and σ_S/R_f are the mean and standard deviation of the random variable $m S_T$). The quantity δ is a “distance to default,” being the number of standard deviations (of $m S_T$) by which the value of the surplus (being the mean of $m S_T$) exceeds zero. We then have expressions for ζ in terms of these standardised quantities:

$$\begin{aligned}\zeta &= \sigma' \phi\left(\frac{\mu'}{\sigma'}\right) - \mu' \left(1 - \Phi\left(\frac{\mu'}{\sigma'}\right)\right), \\ &= \sigma' (\phi(\delta) - \delta (1 - \Phi(\delta))).\end{aligned}$$

The behaviour of the credit discount can be inferred by graphing these quantities. Figure 1 shows the behaviour of the credit discount ratio (scaled by σ') as a function of distance to default. This quantity is unbounded as δ approaches $-\infty$, indicating that the credit discount can exceed one. This is not an error, and corresponds to (extreme) cases where the value of the negative part of the creditors' payoff outweighs that of the positive part, so their claim is worth a negative amount. The firm would need to be more strongly capitalised in order to be able to do business with informed creditors in such cases. Figure 2 shows the credit discount as a function of the standard deviation of the surplus, with the value of the surplus (expressed as a multiple of the value of the claims) held constant. When the surplus is negative, a discount applies even if the standard deviation of the surplus is zero (as the assets are worth less than the claims — the firm is insolvent).

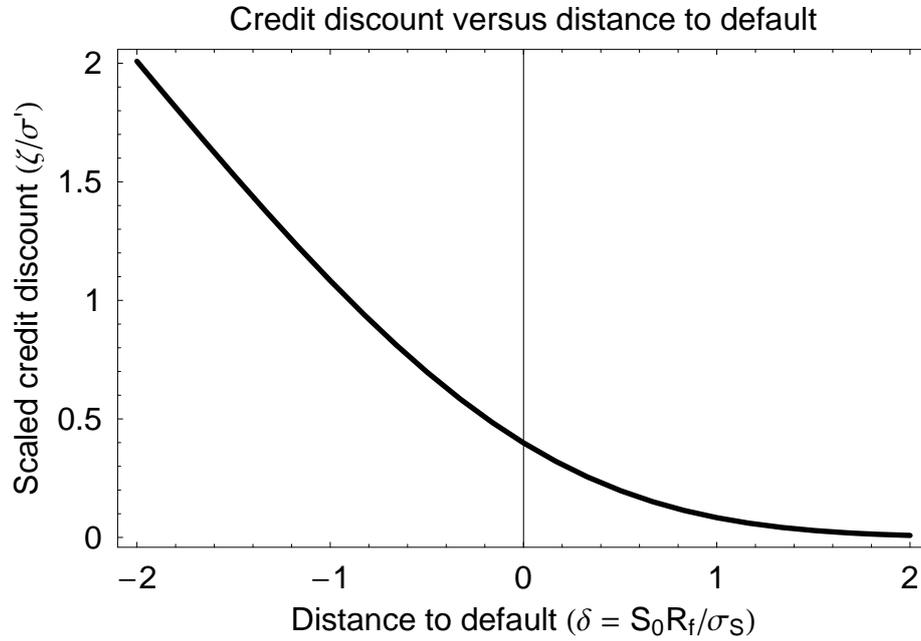


Figure 1: Credit discount per unit of scaled surplus standard deviation ($\frac{\zeta}{\sigma_r}$) versus distance to default (δ).

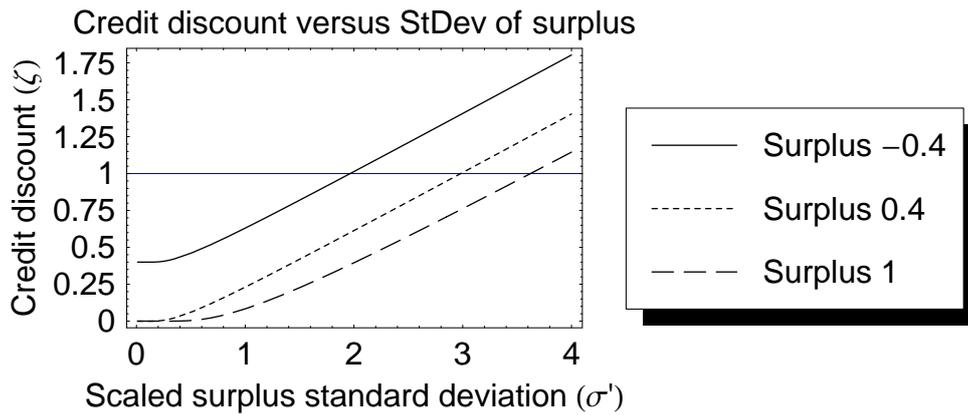


Figure 2: Credit discount versus the scaled standard deviation of the surplus (σ').

For each curve we can see the point where the credit discount first exceeds one, so that the claims, in aggregate, become worthless. The more realistic parts of the curves are those where the credit discount ratio is small.

9 Conclusions

We have introduced the notion of a marginal stochastic discount factor, and shown how one may derive the marginal stochastic discount factor for the exponential-utility capital asset pricing model. The advantage of this formulation of the CAPM is that it allows us to obtain equilibrium prices for assets which are derivatives of CAPM assets, even though these derivatives do not satisfy the CAPM. We have illustrated the practical application of this discount factor by showing how the value of the liabilities of a firm with a normal surplus may be computed. In a subsequent paper we will expand upon this example to value the balance sheet components of archetypal firms — debt- and equity-funded corporations, insurance companies and banks — and examine the properties of risk-based capital allocation schemes in such firms.

We have arrived at the stochastic discount factor for the CAPM by setting out the work of Bühlmann and Wang in the stochastic discount factor framework. In the complete-markets Bühlmann economy, where agents have exponential utility functions, the stochastic discount factor is an exponential function of aggregate consumption. The stochastic discount factor is thus strictly positive, so this economy is arbitrage-free. In the Wang-Bühlmann economy, which is a specialisation of the Bühlmann economy where aggregate consumption is normally distributed, the stochastic discount factor is thus log-normally distributed. Within the Wang-Bühlmann economy, we have identified two classes of assets for which we can obtain an explicit form for the marginal discount factor. The Wang assets are those which are normal-copula with consumption. The normal assets are a subset of the Wang assets, being those which are jointly-normal with consumption. These assets form a vector subspace of the space of all assets. They satisfy the Capital Asset Pricing Model (CAPM).

References

- [Büh80a] Hans Bühlmann. An economic premium principle. *ASTIN Bulletin*, 11:52–60, 1980.

- [Büh80b] Hans Bühlmann. The general economic premium principle. *ASTIN Bulletin*, 14(1):13–21, 1980.
- [Coc01] John H. Cochrane. *Asset Pricing*. Princeton University Press, Princeton, New Jersey, 2001.
- [Joh03] Mark E. Johnston. Re-visiting the principles of insurance pricing using modern economic valuation methods. Submitted to the Journal of the Institute of Actuaries of Australia, 2003.
- [LW01] Stephen F. LeRoy and Jan Werner. *Principles of Financial Economics*. Cambridge University Press, Cambridge, United Kingdom, 2001.
- [P+98] Harry H. Panjer et al. *Financial Economics: with applications to investments, insurance and pensions*. The Actuarial Foundation, Schaumburg, Illinois, USA, 1998.
- [She04] Michael Sherris. Solvency, capital allocation and fair rate of return in insurance. Preprint, University of New South Wales, 2004.
- [Ven03] Gary G. Venter. Tails of copulas. Casualty Actuarial Society, 2003.
- [Wan02] Shaun S. Wang. A universal framework for pricing financial and insurance risks. *ASTIN Bulletin*, 32(2):213–234, 2002.
- [Wan03] Shaun S. Wang. Equilibrium pricing transforms: New results using Bühlmann’s 1980 economic model. To appear in *ASTIN Bulletin*, 2003.