
H Asset pricing

‘The centrality of the trade-off between risk and expected return infuses all investment decisions’. Bernstein (2007, p. xvii).

H.1 The consumption based model

The starting point for any theory of asset pricing must be the investor’s first-order condition for savings and portfolio choice – the basic consumption based model.¹ The consumer has a utility function defined over current and future consumption:

$$U(c_j, c_{j+1}) = u(c_j) + \gamma E_j[u(c_{j+1})]$$

An investor must decide how much to save and how much to consume. Let an asset have an uncertain payoff of x_{j+1} in time $j+1$. The investor can buy and sell as much of the payoff x_{j+1} as he wishes, at price p_j . The marginal utility loss of consuming a little less today and buying a little more of the asset should equal the marginal utility gain of consuming a little more of the asset’s payoff in the future. That is, the consumer’s first order condition is:

$$p_j = E(m_{j+1}x_{j+1})$$
$$m_{j+1} = \gamma \frac{u'(c_{j+1})}{u'(c_j)}$$

where m_{j+1} is the stochastic discount factor or marginal rate of substitution and E is the expected value, based on information at time j . If the equation does not hold, then the investor would buy more or less of the asset (and change current and future consumption) until it does hold.

The equation can be rewritten as

$$p_j = \frac{E(x_{j+1})}{R_{j+1}^f} + \text{cov}(m_{j+1}, x_{j+1}) = \frac{E(x_{j+1})}{R_{j+1}^f} + \frac{\text{cov}[\gamma u'(c_{j+1}), x_{j+1}]}{u'(c_j)}$$

¹ This section is based on Cochrane (2005, pp. 3-17, 35-37).

where the risk free gross return is $1+r_{j+1}^f = R_{j+1}^f = 1/E(m_{j+1})$. That is, expected returns are proportional to the covariance of returns with the discount factor. The first term on the right hand side is the asset's price in a risk neutral world. The second term is a risk adjustment. An asset whose payoff covaries negatively with the discount factor has its price decreased. Marginal utility declines as consumption increases, so the asset's price decreases if its payoff covaries positively with consumption. Such an asset makes the consumption stream more volatile. If the asset covaries negatively with consumption, it has a high price.

That is, it is the covariance of a payoff with consumption rather than its variance that determines its riskiness to an investor. What matters to the investor is the volatility of his consumption.

If the asset price is set equal to 1, then the payoff is a gross return – say R^k for asset k . Then the equation can be written as:

$$E(R^k) - R^f = -R^f \operatorname{cov}(m, R^k) = -\frac{\operatorname{cov}[u(c_{j+1}), R_{j+1}^k]}{E[u'(c_{j+1})]}$$

That is, all assets have an expected return equal to the risk free rate plus a risk adjustment. Assets whose returns covary positively with consumption make consumption more volatile and must offer a higher expected return to induce investors to hold them. Idiosyncratic risk, uncorrelated with the discount factor, does not affect asset prices or generate a higher return. Only systematic, or market, risk generates a risk correction in the price, or a return different from the riskless return (a risk premium in the discount rate).

The basic pricing equation can be written as:

$$E(R^k) = R^f + \left(\frac{\operatorname{cov}(R^k, m)}{\operatorname{var}(m)} \right) \left(-\frac{\operatorname{var}(m)}{E(m)} \right) = R^f + \beta_{k,m} \lambda_m$$

where $\beta_{j,m}$ is the regression coefficient of the return R^k on m (which varies from asset to asset). λ_m is often interpreted as the price of risk and $\beta_{k,m}$ as the quantity of risk in each asset. The price of risk depends on the volatility of the discount factor. This is a beta pricing model, where each expected return is proportional to $\beta_{k,m}$.

So far the only assumption made is that investors can make small marginal investments in the asset. The basic pricing equation does not assume complete markets, a representative investor, normally distributed asset returns, two periods, quadratic utility, no labour income or market equilibrium. The equation applies to each individual, for each asset to which he has access and for any two periods of a multi-period model. Further, the assumption that the consumer maximises time and

date separable expected utility could be dropped. The basic consumption pricing equation would hold for any general non-separable utility function. The relation between the discount factor and real variables would, however, be more complicated. It would be a function of many variables, not just the discount on future utility and present and future consumption.²

H.2 The Capital Asset Pricing Model

Further assumptions can be added to derive other asset pricing models, such as the widely used Capital Asset Pricing Model (CAPM). The CAPM is an example of a factor pricing model, where the discount factor m_{j+1} is a linear function of factors. The factors are plausible proxies for marginal utility. That is, factor models are special cases of the consumption based model – special cases that allow marginal utility growth to be proxied from some other variable.³

In the CAPM the market portfolio is the only source of the consumer's income so that $m_{j+1} = a + bR_{j+1}^W$, where R^W is the return on a claim to total wealth, often proxied by a broad share market portfolio. a and b are parameters. Future consumption is funded solely from returns to portfolios of securities. For example, there is no labour income.

The CAPM's simplifying assumptions make asset prices linear in the single risk factor, the market return. For example, that security returns are jointly normally distributed. Alternatively that investors have quadratic utility. In a multi-period model the assumptions that returns are independent and identically distributed (no shifts in the investment opportunity set), and the interest rate and relative commodity prices are constant, need to be added. If consumption risk changes over time, the CAPM fails to hold.⁴

Further, investors have homogenous expectations and there are no transactions costs. A risk free security exists and there are no borrowing constraints.⁵ Alternatively, in the absence of a riskless asset, that there are no short-selling constraints so that the investor can form a zero beta security.

² Cochrane (2005, pp. 35-36).

³ Cochrane (2005, pp. 149-52).

⁴ See Cochrane (2005, pp. 152-55); Jones (2008, p. 129).

⁵ See Jones (2008, pp. 123-29). These assumptions mean all consumers face the same capital market line, which relates portfolio variance and return.

The CAPM prices assets relative to a market portfolio. The CAPM uses the return on the market portfolio as the single factor. It can be expressed as a beta pricing model:

$$E(R^k) = R^f + \beta_{kM}[E(R^M) - R^f] = \beta_{kM}E(R^M) + (1 - \beta_{kM})R^f$$

where $\beta_{kM} = \sigma_{kM} / \sigma_M^2 = COV(R^M, R^k) / VAR(R^M) = \rho_{kM} \sigma_k / \sigma_M$ measures the amount of market risk in security j , the risk in the asset that cannot be diversified away and $\rho_{kM} = \sigma_{kM} / \sigma_k \sigma_M$ is the coefficient of correlation of the return on security k and the market return, R_M . This CAPM asset pricing equation is known as the security market line.

In the CAPM, all investors measure and price risk identically. Asset prices depend solely on the expected cash payoff, the quantity of market risk provided by the asset, the risk free rate and the price of risk — which are market-determined variables. Individuals who perceive the same distribution of payoffs for a risky asset will price it exactly the same way, independent of their individual utility functions and attitudes to risk.⁶ The CAPM approach does not require everyone has the same utility function — just common expectations

In the CAPM, what matters is not the variability of an asset's income flows themselves, but the contribution they make to overall portfolio risk (and, ultimately, to consumption risk). For example, holding a diversified portfolio can eliminate diversifiable risk, unexpected losses on some assets are offset by gains on others.

Market risk, by contrast, is associated with economy-wide increases or decreases in asset values (resulting in changes in aggregate consumption), so it cannot be eliminated through portfolio diversification. Market risk arises from the volatility of the economy and from associated changes in the value of aggregate wealth. Because those changes create undesired uncertainty, they are costly to investors and are reflected in market prices.

Only non-diversifiable, or market, risk is costly. The level of market risk associated with an asset determines the return demanded by the market above the risk free return, that is the risk premium. Assets that only have diversifiable risk earn the riskless rate of return (a zero risk premium). Assets which are positively correlated with market return (contain market risk) must earn a higher rate of return. The higher the correlation with market return, the higher the required rate of return. Assets that are negatively correlated with the market reduce portfolio risk (they provide insurance) and earn less than the riskless return.

⁶ Copeland and Weston (1983, p. 196).

The risk premium is about how much market risk is contained in the stream of payments, it is not about downside risk, which should be accounted for in the expected value estimates.

H.3 Multifactor models

The CAPM has the advantage of convenience and ‘seductive simplicity’.⁷ Although the CAPM is widely used and has proved to be a powerful tool for assessing risks, it fails many empirical tests: much variation in expected returns is unrelated to market beta. That is not surprising, as the CAPM is based on many unrealistic assumptions, such as that investors only care about the mean and variance of one-period portfolio returns, which miss important dimensions of risk. For example, investors also care about how their portfolio return covaries with labour income and future investment opportunities (such as changing expectations of future returns) — which affect marginal utility. Market beta, then, is not a complete description of an asset’s risk and so does not explain all differences in expected asset returns.⁸ There are assets, portfolios and strategies whose average returns cannot be explained by their market beta. Additional sources of priced risk beyond movements in the market portfolio are needed to explain why some average returns are higher than others.⁹

An extension has been multifactor models, such as the intertemporal capital asset pricing model (ICAPM), where investors care how their wealth varies with future state variables that affect their opportunities to invest or consume the asset’s payoff, such as labour income, the prices of consumption goods, investment opportunities and financial distress factors. Multifactor models add additional beta, along with a market beta, to explain expected returns.

Multifactor models measure risk by covariances with several common factors — usually based on empirical fit rather than theory (which gives rise to the data-snooping critique), arguing that the factors reflect unidentified state variables that produce undiversifiable risks (covariances) in returns that are not captured by the market return and are priced separately from market betas.¹⁰

Multifactor models dominate the empirical description, performance attribution and explanation of average returns, but give little guidance about choosing factors or answering what forces determine factor risk prices. What are the sources of

⁷ Fama and French (2004, p. 44).

⁸ See Fama and French (2004).

⁹ Cochrane (2005, pp. 390-91).

¹⁰ See for example, Fama and French (2004).

macroeconomic risk that drive asset prices? It seems a recession or financial distress factor lies behind many asset prices. People forgo large return premiums to avoid assets that fall in recessions. For example, it takes a risk higher premium to get people to hold stocks at the bottom of a recession. Asset returns depend on standard market beta and four liquidity premiums: expected returns are higher for stocks that are illiquid on average, for stocks whose illiquidity gets worse when the market goes down, for stocks whose price goes down when the market gets more illiquid, and for stocks which become more illiquid when the market gets more illiquid.¹¹

H.4 Asset return puzzles

The equity premium puzzle

The equity premium puzzle is that the standard consumption-based model of asset pricing can only explain the observed equity premium by assuming unrealistic levels of risk aversion.¹²

The Sharpe ratio, a measure of the ratio of reward to risk in asset markets, is:

$\frac{E(R^k) - R^f}{\sigma_k}$ for asset k (it is proportional to the price of risk derived in H.1 and is sometimes described as the market price per unit of risk). It is invariant to leveraging.

From H.1, the basic asset pricing equation in terms of an asset's return is:

$$1 = E(mR^k) = E(m)E(R^k) + \rho_{m,R^k} \sigma \sigma_k$$

Where $\sigma = \sqrt{\text{var}(m)}$ is the standard deviation of the stochastic discount factor, σ_k is the standard deviation of R^k and ρ_{m,R^k} is a correlation co-efficient. This can be written as: $E(R^k) = R^f - \rho_{m,R^k} \frac{\sigma}{E(m)} \sigma_k$. As the absolute value of correlation co-

efficients cannot be greater than 1, this gives: $\left| \frac{E(R^k) - R^f}{\sigma_k} \right| \leq \frac{\sigma}{E(m)}$

¹¹ Cochrane (2004, p. 9).

¹² This section summarises Cochrane (2005, pp. 17-22, 455-59).

Now for power utility, $m = \gamma \left(\frac{c_{t+1}}{c_t} \right)^{-\eta}$ (see appendix F). If consumption growth is normally distributed, then we get the Hansen-Jagannathan bound:

$$\left| \frac{E(R^k) - R^f}{\sigma_k} \right| \leq \frac{\sigma}{E(m)} = R^f \sigma \approx \eta \sigma_c \quad \text{where } \sigma_c \text{ is the standard deviation of consumption growth.}$$

In the United States, historic measures of the equity premium for the broad share market are around 6 to 8 per cent and the standard deviation of the market return around 16 to 20 per cent, giving a Sharpe ratio of 0.3 to 0.5.¹³

Complete markets mean that individuals can diversify any idiosyncratic risk in consumption, so that only fluctuations in aggregate consumption are relevant for consumers. The standard deviation of aggregate consumption, σ_c , is about 1 to 2 percent. The Hansen-Jagannathan bound implies that the co-efficient of relative risk aversion (η) must be at least 15 (for $\sigma_c = 2$ and a Sharpe ratio of 0.3) and could be as high as 50 (for $\sigma_c = 1$ and a Sharpe ratio of 0.5), an incredibly high level of risk aversion.

Many believe that that the historically high risk premium reflects luck and the expected future risk premium is in the range 3 to 5 per cent, giving a Sharpe ratio of 0.15 to 0.3 (the ratio in Australia for the past 38 years was 0.2).¹⁴ That reduces the required co-efficient of risk aversion to as low as 7.5. But even that is difficult to explain. Empirical estimates of η range from 0.2 to 4 (with most between 1 and 2).¹⁵ A η between 1 and 2 and a standard deviation of consumption growth of between 1 and 2 per cent imply a Sharpe ratio between 0.01 and 0.04 and, at most, a risk premium of 0.008 (less than 1 per cent). Observed levels of these variables are greater by a factor of 10.

Ways of generalising the calculation make matters worse. Even broad share market indexes are only a portion of the true market portfolio. But that means they likely lie inside the mean-variance frontier and so their Sharpe ratio understates the true Sharpe ratio. The above calculations assume that asset returns and consumption growth are perfectly correlated. In the United States aggregate consumption has a 0.2 correlation with the market return, raising the required level of risk aversion fivefold (to 75–250).¹⁶

¹³ Cochrane (2005, p. 456); Mehra and Prescott (2003, p. 908).

¹⁴ See appendix G.

¹⁵ Zhuang et al. (2007), Weitzman (2007, p. 709).

¹⁶ Cochrane (2005, p. 21).

The risk free rate puzzle

If the level of risk aversion was really as high as the simple consumption based model suggests, then a risk free rate puzzle would arise. For example, individuals should strongly want to smooth consumption over time, because consumption shortfalls have a much larger effect on utility than consumption increases. As the economy is growing, individuals should be keen to borrow to increase current consumption. Everyone trying to borrow should result in high interest rates, but the observed riskless rate is low. In a certain world, the riskless rate is $r^f = \theta + \eta g$.¹⁷ g (average consumption growth) is around 2 percent, and so an η equal to 15 or 50 implies a riskless interest rate of over 30 or 100 percent, so long as θ (the pure rate of time preference) is not negative (that is, so long as there is a preference for current over future utility). These rates are vastly above observed rates.

The Ramsey formula for the risk free rate can be extended to account for uncertainty. If consumption growth is independently and identically distributed normal with a known mean g and variance σ_c^2 , then the formula for the risk free interest rate is:¹⁸

$$r^f = \theta + \eta g - \frac{\eta^2 \sigma_c^2}{2}$$

In a world of uncertainty, the real interest rate is also negatively related to the degree of risk aversion (for example, there is a precautionary motive for saving, which decreases interest rates). For example, extremely high levels of risk aversion can produce low risk free rates, even puzzlingly low. For example, if we substitute $g = 2$ percent, $\theta = 2$ percent and $\eta = 100, 200$ into the formula for the risk free rate under uncertainty, we get a risk free rate of 2 percent and -398 percent.

Although high levels of risk aversion can (almost by co-incidence) give reasonable levels of the risk free rate, small changes in parameters would result in dramatic changes in the risk free rate. That is why such high levels of risk aversion are inconsistent with real world observations. A high η means that consumers are unwilling to substitute consumption over time, so it would take large variations in the interest rate to get them to adapt to small variations in consumption growth. For example, a one percentage point increase in consumption growth would raise interest rates by η percentage points, totally implausible if η is high.¹⁹ Observed real

¹⁷ This formula is derived in appendix F.

¹⁸ This model is set out in Weitzman (2007, p. 710). The equations are derived in Cochrane (2005, p. 12).

¹⁹ See Cochrane (2005, pp. 457-458).

interest rates are not only low, but relatively stable (for example, annual realised real returns on 90 day Treasury Bills since 1971 had a standard deviation of 3.5 per cent).

Reasonable parameter values, such as $\theta = 0$ to 2 per cent, $\eta = 1$ to 2, $g = 1$ to 2 per cent and $\sigma_c = 1$ to 2 per cent, give a negligible risk premium and a risk-free rate of anywhere from 1 to 6 per cent. In the US the return on equities has averaged around 7 per cent real and on short term Treasury bonds around 1 per cent, and the puzzle is the gap between them. The Ramsey equation does not resolve the puzzle of which rate to choose.

Explanations of the equity premium puzzle

A high equity premium must come from a lot of risk or a lot of risk aversion. The explanations of the equity premium puzzle fall into three main categories: preference modifications, individual consumption risk and fear of disasters.²⁰ No one explanation has been universally accepted, but each explains part of the premium, and together, perhaps all.

Preference modifications

Modifying preferences can make consumers extremely adverse to share market losses and yet also imply a low risk free rate. For example, in the standard approach the one parameter determines both risk aversion and time preference (η is the coefficient of risk aversion and the reciprocal of the elasticity of intertemporal substitution). Some authors have suggested utility functions (Epstein-Zin preferences) which break this link, and so allow high risk aversion without the risk free rate puzzle.

Standard preferences assume time separability and state independence. Changes to these assumptions that make consumers effectively highly risk averse include making utility depend on past own consumption (habits), past consumption of others or the current consumption of others, so that changes in aggregate growth have a large effect on well-being.²¹

Another approach are ad hoc behavioural explanations. For example, Siegal and Thaler (1997) put the equity premium puzzle down to loss aversion on the part of investors combined with annual portfolio evaluations. Others claim investors

²⁰ Survey of the equity premium puzzle include DeLong and Magin (2009); Mehra and Prescott (2003); Mehra (2006, 2003); Cochrane (1997) and Kotcherlakota (1996).

²¹ See Jones (2008, pp. 147-48); Cochrane (2005, pp. 465-81; 1997, pp. 19-22).

irrationally over-estimate the riskiness of equity because of behavioural biases, such as focusing on the short-term.

These preferences can explain at least some of the equity premium and risk free rate puzzle. But they still require high levels of risk aversion to account for the equity premium.

To the extent that preference based explanations of the equity premium are correct, the resulting premium should be included in the social discount rate. It represents the real cost of risk to consumers based on their own evaluations.

Individual consumption risk

In the consumption based pricing model, consumers have the same (aggregate) consumption risk because they costly eliminate diversifiable risk, Incomplete markets, borrowing constraints and transactions costs can stop consumers from diversifying risks and make individual consumption risk higher than aggregate consumption risk.²² For example, most people have a job that could be under threat in a recession. Uninsurable human capital risk can make people unwilling to hold the additional aggregate risk associated with shares. It can also encourage saving, driving down the risk free rate.

Other reasons for increased individual consumption risk include consumption commitments and concentrated share ownership.²³ Only a small minority of population that holds significant amounts of shares. Their consumption is more volatile and more highly correlated with the stock market return. Consumption commitments amplify risk aversion.

Uninsured individual risk is not an obvious solution to the equity premium puzzle. Individual consumption is not volatile enough to satisfy the Hansen-Jagannathan bound, and is less correlated with stock returns than aggregate consumption – and the more volatile it is, the less correlated. Idiosyncratic shocks uncorrelated with asset returns have no effect on asset pricing. Shocks correlated with asset returns are quickly traded away.²⁴ For example, by using borrowing to create ‘home made consumption smoothing’. Although borrowing constraints may limit consumers’ ability to smooth, in that case their discount rates would be high.

Consantinides and Duffie do present an idiosyncratic risk model, but prevent people self-insuring by making the shocks permanent. The variance of idiosyncratic risk

²² See, for example, Mankiw (1986).

²³ See Mankiw and Zeldes (1991) and Chetty and Szeidl (2007).

²⁴ Cochrane (2005, p. 459-60; 1997, pp. 23-24).

rises when the market declines. The model shows how high consumption volatility is consistent with low aggregate volatility. But the distribution of consumption does not spread out as much as the model requires to explain all the equity premium. With any plausible characterization of the actual labour income process, the model requires high consumption volatility or high risk aversion to explain the observed risk premium.

To the extent that high individual risk causes the equity premium, it represents a real cost of risk to individuals and should be included in the social discount rate unless the government has some advantage that allows it to overcome the imperfections and diversify risk that individuals cannot. If a government project imposes risk on individuals, the market price of risk indicates the cost of that risk .

Fear of disasters

Weitzman (2007) points out that the standard finance models assume agents know the economy's stable structural parameters: such as the mean and variance of the (assumed) normal distribution of the economy's growth rate. But if instead these structural parameters are uncertain, evolving and have to be discovered, then the major asset return puzzles (including the equity premium puzzle) become explicable because of investor aversion to structural uncertainty and fear of a negative growth 'disaster'. The puzzles involve

skittish investors nervously reacting with unsure expectations to deeper forces of shifting structure ... to an outside observer it looks like inside investors are being rewarded by an inexplicably-high empirical asset return, while actually they are bearing the extra risk of rare disasters in the left tail of the distribution that happen not to have materialized within the limited sample.²⁵

The result is the so-called peso problem: where investors take into account some probability of a rare disaster when forming expectations. If the disaster is not observed in the sample, realised returns seem inexplicably high, but they are compensating investors for the risk of rare disasters that happen not to have materialized.²⁶ To complicate matters, the expected probability of disasters may change over time.

Weitzman concludes that makes sense to use the market premium rather than the one derived from observed variance in aggregate consumption

there is some theoretical justification for treating the subjective variability of the future growth rate as if it were equivalent in welfare to the observed variability of a

²⁵ Weitzman (2007a, p. 1104).

²⁶ Weitzman (2007, p. 1104).

comprehensive economy-wide index of equity-wealth returns. For this as-if- REE high-growth-variability interpretation, the simple standard model of asset pricing may have the potential to be a decent shortcut conceptualization of what is actually happening in a complicated ever-changing world where unforeseen bad events—scary, disruptive, and without precedent— may evolve at any future time.²⁷

That is, the share market is much more volatile than consumption, but it more accurately reflects the true risks to consumption.

Weitzman’s analysis dovetails with the Barro approach. Barro has a series of papers explaining the premium puzzle by a small probability of a disaster.²⁸ It has to be a disaster that reduces the returns to equities by more than bonds — such as the Great Depression.

Barro and Ursua (2008) find that the probability of a ‘consumption disasters’ (a greater than 10 percent decline in aggregate consumption), estimated from historical data, justifies the equity premium. They assume a known proportional reduction in consumption occurs with a known probability. Weitzman emphasises that asset puzzles dissolve when we also account for the process of learning about these parameters.

As the benefits from government investments are likely to vary with aggregate consumption and fall during consumption catastrophes, a reasonable assumption is that government investments face the same risk as private investments and the cost of capital should include the market risk premium caused by fear of disasters.

Other explanations

Transactions costs can explain the equity premium when equity is more costly to trade than debt. The equity premium may reflect the lower liquidity of equity compared with bonds. The risk free asset should be free from liquidity risk, there may be a premium for the risk from holding less liquid assets.²⁹

If the high return on equities is to compensate for the costs that arise because of their lack of liquidity, the bond return still does not correctly measure the costs from government projects. If anything the liquidity premium for government assets should be higher because from the point of view of the taxpayer, a government

²⁷ Weitzman (2007a, p. 1128).

²⁸ See Barro (2006; 2007; 2009), Barro and Ursúa (2008). See also Wachter (2008), who examines a time varying probability of disaster.

²⁹ See Swan (2006) and Levy and Swan (2008).

investment is even less liquid than equity. Unlike equity it cannot be traded or disposed of by an individual taxpayer.

If government does have an advantage in raising money, then that is an argument for expanding government financing role. But it should still invest in the projects with the highest return, an objective promoted by using a market-based discount rate.

Further, the equity premium puzzle is the high Sharpe ratio for stocks. Cochrane points out that

There are large Sharpe ratios *between* stocks (as in the value-growth premium studied by Fama and French) ignoring bonds all together. High sample Sharpe ratios are pervasive in finance and not limited to the difference between stocks and bonds.³⁰

Differences in the liquidity of stocks and bonds cannot cause the entire equity premium.

McGrattan and Prescott (2003) argue that tax and regulatory changes boosted the demand for shares — which raised realised returns but lowered expected returns, accounting for the equity premium puzzle.³¹

H.5 Two approaches to accounting for risk: certainty equivalent v risk premium.

For a one period valuation, using a risk-adjusted discount rate on expected values or the risk free rate on certainty equivalents are alternative ways to account for risk.³² Both should give the same answer. For example, if the CAPM holds, then

$$E(R^k) = R^f + \beta_{kM}[E(R^M) - R^f] = R^f + \psi COV(R^M, R^k) \text{ where}$$

$$\beta_{kM} = COV(R^M, R^k) / VAR(R^M) \text{ and } \psi = [E(R^M) - R^f] / VAR(R^M).$$

Now if P_0 is the price of the asset and P_E is the risky payoff received at the end of the period then $E(R^k) = E(P_E) / P_0 = R^f + \psi COV(R^M, R^k)$.

Rearranging this gives the risk adjusted rate of return valuation formula:

$$P_0 = E(P_E) / [R^f + \psi COV(R^M, R^k)].$$

³⁰ Cochrane (2008, p. 259).

³¹ McGrattan and Prescott (2003).

³² The following is from Copeland and Weston (1983, pp. 195-96).

The value of the asset is the expected value of the risky payment discounted with a risk adjusted discount rate. If the asset has positive market risk, a risk premium is added to the risk free rate. The higher the covariance with the market return, the more market risk it contains and the higher the required rate of return.

An equivalent way to value the asset is to deduct a risk charge from the expected value of the risky payoff to convert it into a certainty equivalent and discount with the risk free return. The certainty equivalent of a risky payment is the smallest certain payment which the household would accept in exchange for the risky one.

The covariance between the risky asset and the market can be written as:

$$\text{cov}(R^M, R^k) = \text{cov}\left[R^M, \frac{P_e}{P_0}\right] = \frac{1}{P_0} \text{COV}(P_e, R^M). \text{ Substituting this into the risk-}$$

adjusted rate of return equation gives:

$$P_0 = \frac{E(P_e)}{R^f + \psi(1/P_0)\text{COV}(P_e, R^M)}. \text{ Rearranging this gives the certainty equivalent}$$

valuation formula:

$$P_0 = \frac{E(P_e) - \psi \text{COV}(P_e, R^M)}{R^f}$$

The risk charge depends on the level of market risk in the benefit flow. When a benefit co-varies with the market return, it contains market risk and the certainty equivalent is less than the expected value. The more market risk, the lower the certainty equivalent.

The CAPM pricing equation can be expressed as expected values discounted by a risk adjusted discount rate or as their certainty equivalents discounted by the risk free rate. The two approaches are equivalent for one-period valuation models and require exactly the same information. They are just different ways of expressing the same calculation.

A major advantage of the standard CAPM model is that the value of an asset depends solely on the expected cash pay offs, the quantity of market risk in the payoffs, the risk free rate and the price of risk — which are all market determined variables. Individuals who perceive the same distribution of payoffs for a risky asset will price it exactly the same way, independent of their preferences and attitudes to risk.³³ The CAPM approach does not require everyone has the same utility function — just common expectations.

³³ Copeland and Weston (1983, p. 196).

If a benefit is uncorrelated with the market (it only has diversifiable risk), its certainty equivalent is equal to its expected value and the appropriate discount rate is the riskless return. If it is negatively correlated with the market, the risk-adjusted discount rate is below the riskless rate, and the certainty equivalent would be greater than the payment's expected value (it provides insurance and reduces portfolio risk).

H.6 Discounting in a multi-period world

When the net benefit flows from a regulation extend over many periods, the conditions under which adding a constant risk premium to the discount rate is a correct way to account for risk are more stringent. It requires the certainty equivalent relative to the expected value of the net benefits to decline steadily over time, which implies the risks of the benefit stream increase over time. That is, $CE_n/(1+r)^n = EV_n/(1+r^*)^n$ where CE_n is the certainty equivalent of a payment received in period n , EV_n is its expected value and r^* is the risk adjusted discount rate. For this to be true, CE/EV must fall systematically with n . It would be correct if the riskiness of income flows is greater the further in the future they are received. That is, if the project's cumulative risk increases at a constant rate so that the certainty equivalent approach requires a larger deduction for risk the further into the future the cash flow is received.

Using the market risk premium is justified when the net benefits from the proposal follow a Weiner process.³⁴ A Weiner process is the same as Brownian motion. The variance of the change in a Weiner process grows linearly with the time horizon. Over the long run, its variance goes to infinity. It is a common assumption in financial economics. The proportional risk associated with holding such an asset grows linearly over time.

The two main reasons why market risk of net cash flows can change over time is investor reassessment of the project risk and changes in aggregate consumption risk.

The standard CAPM can be used in a multi-period model only if the project's beta each period is known with certainty. Intermediate uncertainty is only admissible if it does not affect the beta coefficients.³⁵

³⁴ For example, if $NB_j = E_{j-1}(NB_j)(1 + \varepsilon_j)$ where NB_j is the net benefit at time j , $E_{j-1}(NB_j)$ is its expected value at time $j - 1$ and ε_j is a random variable with zero mean and constant covariance with the return on the market portfolio. Then net benefits at time j evolves as a martingale and the risk-adjusted discount rate can be used to compute its present value. See Copeland and Weston (1983, pp. 361-63).

³⁵ See Brearley and Myers (1991, p. 201) and Jones (2008, pp. 153-156).

When the risk contained in net benefits flows does not increase linearly through the life of the project, using one constant discount rate leads to incorrect decisions, and the certainty equivalent approach is preferable. It is more flexible and allows separate risk adjustment for net benefits received in different periods.

To calculate certainty equivalents is difficult. It is equally difficult to calculate the corresponding discount rate structure, which requires the same information. A common practical solution is to use a constant risk-adjusted discount rate, which uses market information on the current price of market risk.

Myers and Turnbull (1977) find that this traditional CAPM approach gives close to correct answers in a multi-period world providing the right asset beta is used to calculate the discount rate. But that depends on project life, growth trend in expected cash flows and other variables, which are not usually considered important in assessing business risk.³⁶

A constant risk adjusted discount rate gives the correct answer (that is, the same as discounting the certainty equivalents at the riskless rate) if cash flows follow a pure random walk. Otherwise, conventional capital budgeting is wrong, but it often gives close to the right answer. Myers and Turnbull (1997. p. 330) conclude:

no serious errors are introduced by discounting cash flow streams at one-period expected rates of return inferred from observed beta. ... conventional valuation formulas based on discounting expected cash flows give a good approximation to asset values derived from rigorous analysis of equilibrium market values.

But determining beta is more complicated. Beta depends on the link between cash flow forecast errors and forecast errors on the market return. It depends on asset life, the growth trend in cash flows and on the pattern of expected cash flows over time (which affects the relative weights on the components of cash flows).

Further, the discount rate in a multi-period world may include extra risk premiums for other risks — such as a changing riskless rate of return and changes in the economy.

³⁶ Appendix I, section I.5 sets out how beta in the CAPM falls with asset life.