Asset Pricing in Macroeconomic Models*

Paul Söderlind[†]

Stockholm School of Economics

6 September 2002

Abstract

Analysis of financial prices in macroeconomic models rests on two building blocks: the consumption-based asset pricing model and the structure of payoffs. This chapter studies how different modelling choices affect yield curves (real and nominal), risk premia on equity (levered or not), and options. The emphasis is on surveying existing models and to bring out the basic mechanisms and intuition. All results are based on simple stylized facts and analytical pricing expressions.

Keywords: consumption-based asset pricing model, equity premium puzzle, riskfree rate puzzle, Sharpe ratio, cross-sectional variation in returns.

JEL Classification Numbers: G12, E130, E320

---

*Prepared for *Elements of Dynamic Macroeconomic Analysis*, edited by Altug, Chadha, and Nolan.

[†] Stockholm School of Economics and CEPR. *Address:* Stockholm School of Economics, PO Box 6501, SE-113 83 Stockholm, Sweden. *E-mail:* Paul.Soderlind@hhs.se. I thank the editors and two anonymous referees for comments.

1 Introduction

This chapter studies the two main financial building blocks of simulation models: the consumption-based asset pricing model and the definition of asset classes. The aim is to discuss what different modelling choices imply for asset prices. For instance, what is the effect of different utility functions, investment technologies, monetary policies, and leverage on risk premia and yield curves?

The emphasis is on surveying existing models and discussing the main mechanisms behind the results. I therefore choose to work with stylized facts and simple analytical pricing expressions. There are no simulations or advanced econometrics in this chapter.¹ The following two examples should give the flavour. First, I use simple pricing expressions and scatter plots to show that the consumption based asset pricing model cannot explain the cross-sectional variation of Sharpe ratios. Second, I discuss how the slope of the real yield curve is driven by the autocorrelation of consumption by studying explicit log-linear pricing formulas of just two assets: a one-period bond and a one-period forward contract.

The plan of the chapter is as follows. Section 2 deals with the consumption-based asset pricing model. It studies if the model is compatible with historical consumption and asset data. From a modelling perspective, the implicit question is: if my model could create a realistic consumption process and has defined realistic assets (for instance, levered equity), would it then predict reasonable asset returns? Section 3 deals with how assets are defined in simulation models and what that implies for pricing. I discuss yield curves (real and nominal), claims on consumption (one-period and multi-period), options, and levered equity. Section 4 summarises the main findings. Technical details are found in a number of appendices at the end.

2 Problems with the Consumption-Based Asset Pricing Model

This part of the chapter takes a hard look at the consumption-based asset pricing model since it is one of the building blocks in general equilibrium models. The approach is to derive simple analytical pricing expressions and to study stylized facts—with the aim of conveying the intuition for the results.

---

¹See Rouwenhorst (1995) for a survey of simulation evidence; Ferson (1995) and Reffett and Schorfheide (2000) for surveys of careful econometric modelling.

The first sections below look at earlier findings on the equity premium puzzle and the riskfree rate puzzle and study if they are stable across different samples (see the surveys of Bossaert (2002), Campbell (2001), Cochrane (2001), and Smith and Wickens (2002)). The following sections present new evidence on the ability of the consumption-based model to account for the predictability of returns and the cross-sectional variation (across portfolios formed on industries, firm size, dividend-price ratio, book-to-market ratio, and country). The last section discusses if alternative models of the stochastic discount factor (including Epstein-Zin utility, habit persistence, and idiosyncratic risk) can help solving the problems of the consumption-based model.

2.1 The Asset Pricing Equation

This section sets the stage for studying whether the consumption-based asset pricing model can fit the data. The basic pricing equation is specified and some stylized facts are presented.

The basic asset pricing equation says

$$E_{t-1}(R_t M_t) = 1, \quad (1)$$

where R_t is the gross return of holding an asset from period $t - 1$ to t , M_t is a stochastic discount factor (SDF). E_{t-1} denotes the expectations conditional on the information in period $t - 1$, that is, when the investment decision is made. This equation holds for any asset that are freely traded without transaction costs (or taxes), even if markets are incomplete.²

It is convenient to rewrite (1) as³

$$E_{t-1}(R_t) = 1/E_{t-1}(M_t) - \text{Cov}_{t-1}(R_t, M_t)/E_{t-1}(M_t), \quad (2)$$

which expresses the conditional expectations of the return in terms of the conditional expectation of the SDF and the conditional covariance of the return and the SDF.⁴

---

²The existence of such an SDF is guaranteed if the law of one price holds (portfolios with the same payoffs have the same prices). The SDF is unique if, in addition, markets are complete (all risk can be insured). See, Harrison and Kreps (1979), Ross (1978), and Hansen and Richard (1987); for textbook treatments see LeRoy and Werner (2001) and Cochrane (2001).

³ $\text{Cov}(x, y) = E(xy) - E(x)E(y)$.

⁴A conditional variance, $\text{Var}_{t-1}(x_t)$, is the variance of the forecast error $x_t - E_{t-1}(x_t)$; a conditional covariance, $\text{Cov}_{t-1}(x_t, y_t)$, is the covariance of the forecast errors $x_t - E_{t-1}(x_t)$ and $y_t - E_{t-1}(y_t)$.

Two applications of this equations will be particularly useful in the subsequent analysis. First, a (gross) riskfree rate, R_{ft} , has no conditional covariance with the stochastic discount factor (but it may still be a random variable), so (2) implies that the expected riskfree rate is the inverse of the expected SDF

$$E_{t-1}(R_{ft}) = 1/E_{t-1}(M_t). \quad (3)$$

Second, consider the gross returns on two assets (one could be a riskfree asset) and let R^e be the excess return, that is, the difference between the two returns. Apply (2) on both returns and take the difference to get

$$E_{t-1}(R_t^e) = -\text{Cov}_{t-1}(R_t^e, M_t)/E_{t-1}(M_t), \quad (4)$$

which shows that a positive risk premium requires a negative correlation with the SDF.

In a consumption-based model, (1) is the Euler equation for optimal saving in $t - 1$ where M_t is the ratio of marginal utilities in t and $t - 1$, $M_t = \beta u'(C_t)/u'(C_{t-1})$.⁵ I will focus on the case where the marginal utility of consumption is a function of consumption only, which is by far the most common formulation. This allows for other terms in the utility function, for instance, leisure and real money balances, but they have to be additively separable from the consumption term. With constant relative risk aversion (CRRA) γ , the stochastic discount factor is

$$M_t = \beta(C_t/C_{t-1})^{-\gamma}, \text{ so} \quad (5)$$

$$\ln(M_t) = \ln(\beta) - \gamma \Delta c_t, \text{ where } \Delta c_t = \ln(C_t/C_{t-1}). \quad (6)$$

The second line is only there to introduce the convenient notation Δc_t for the consumption growth rate.

The next few sections study if the pricing model consisting of (1) and (5) can fit historical data. To be clear about what this entails, note the following. First, general equilibrium considerations will not play any role in the analysis: the production side will not be even mentioned. Instead, the focus is on one of the building blocks of an otherwise unspecified model. Second, complete markets are not assumed. The key assumption is rather that the basic asset pricing equation (1) holds for the assets I analyse (broad US/international stock/bond portfolios). This means that the representative US investor can trade in these

---

⁵See Breeden (1979) and Lucas (1978) for early analyses of the consumption-based model.

assets without transaction costs and taxes (clearly an approximation). Third, the properties of historical (ex post) data are assumed to be good approximations of what investors expected. In practice, this assumes both rational expectations and that the sample is large enough for the estimators (of various moments) to be precise.

My approach is to derive simple expressions which highlights the basic mechanisms and makes the analysis more easily adaptable to other data sets (sample periods/countries/assets). These simple expressions are obtained by making assumptions about the distribution of consumption and returns.

As an example of the pros and cons of this approach, suppose we want to analyse how different values of the risk aversion γ affects the unconditional mean of the SDF, $E(M_t)$. If we make no assumptions about the distribution of consumption, then we must essentially take a data set on consumption growth, construct the variable on the right hand side of (5) for a value of γ and estimate its mean. We then repeat this for other values of γ and then describe (for instance, by plotting) how the mean varies with γ . This approach is cumbersome and the findings are hard to generalize. In contrast, by assuming that consumption growth is normally distributed, then we can express $E(M_t)$ as a simple function of γ and the mean and standard deviation of consumption growth.⁶ This second approach makes it much easier to interpret and extend the findings, although it comes at the price of only getting approximate answers (since data is probably not exactly normally distributed).

Table 1 shows the key statistics for quarterly US real returns and consumption growth for the sample periods 1947–2001 and 1970–2001. The means and standard deviations are expressed in percent, and have been annualized by multiplying the mean by four and the standard deviation by two (as would be appropriate for an annual sum of a quarterly iid variable). Two sample periods are used since it is of interest to see how stable the results are, but also because some of the international data used later in the section is only available from 1970.

The table uses a star (*) to indicate which numbers that are significantly different from zero at the 5% significance level. The tests are based on estimating the moments by GMM, using a Newey and West (1987) estimator for the sampling uncertainty of the moments, and then applying the delta-method (see Appendix B). This means that the tests are only valid asymptotically—but are consistent even if data is not iid (as assumed by

---

⁶The result is $E(M_t) = \beta \exp[-\gamma E(\Delta c_t) + \gamma^2 \sigma(\Delta c_t)^2 / 2]$.

	Mean	Std	Sharpe ratio	Auto- corr	Correlation with	
					Δc_t	Δc_{t+1}
<u>1947–2001</u>						
Consumption growth, Δc_t	1.94*	1.10*		0.20*		0.20*
Stock market excess return	7.93*	15.98*	0.50*	0.05	0.14*	0.22*
Long gov bond excess return	1.15	9.38*	0.12	−0.04	−0.10	0.08
1-month T-bill	1.00*	1.57*		0.58*	0.13	0.27*
<u>1970–2001</u>						
Consumption growth, Δc_t	1.89*	0.91*		0.41*		0.41*
Stock market excess return	6.23*	17.87*	0.35*	0.01	0.16*	0.32*
Long gov bond excess return	3.06	11.43*	0.27	−0.08	−0.10	0.21*
1-month T-bill	1.56*	1.44*		0.69*	0.10	0.23*

Table 1: US quarterly real returns and consumption growth (annualized, %). This table shows summary statistics of US real returns and consumption growth rate for 1947Q2–2001Q4 and 1970Q1–2001Q4. A star (*) denotes significance at the 5% level, based on GMM/delta method inference, using a Newey-West estimator with one lag (see Appendix B). Consumption is real per capita consumption of services and non-durable goods. The consumption growth rate in quarter t is $\Delta c_t = \ln(C_t/C_{t-1})$, where C_t is the consumption level in quarter t . Returns are real returns. Excess returns are real returns in excess of the real return on a one-month Treasury Bill. To annualize, quarterly means are multiplied by 4 and quarterly standard deviations by 2, as appropriate if the variable is a random walk. See Appendix A for details on data sources and transformations.

standard tests).

We see, among other things, that consumption has a standard deviation of only 1% (annualized), the stock market has had an average excess return (over a T-bill) of 6–8% (annualized), and that returns are only weakly correlated with consumption growth. These figures will be important in the following sections. Two correlations with consumption growth are shown, since it is unclear if returns should be related to what is recorded as consumption this quarter or the next. The reason is that consumption is measured as a flow during the quarter, while returns are measured at the end of the quarter.

2.2 The Equity Premium Puzzle

This section studies if the consumption-based asset pricing model can explain the historical risk premium on the US stock market.

To discuss the historical average excess returns, it is convenient to work with the

unconditional pricing equation $E(R_t M_t) = 1$, which follows from applying the law of iterated expectations on (1). We can then derive expressions like (2)–(4) for unconditional moments instead.

To highlight the basic problem with the consumption-based model and to simplify the exposition, I assume that the excess return, R_t^e , and the log SDF, $\ln(M_t)$, have a bivariate normal distribution.⁷ Stein’s lemma⁸ then gives

$$\text{Cov}(R_t^e, M_t) = \text{Cov}[R_t^e, \ln(M_t)] \times E(M_t). \quad (7)$$

We use this in the unconditional version of (4) to get the excess return and Sharpe ratio as

$$E(R_t^e) = -\text{Cov}[R_t^e, \ln(M_t)], \text{ and} \quad (8)$$

$$E(R_t^e)/\sigma(R_t^e) = -\rho[R_t^e, \ln(M_t)] \times \sigma[\ln(M_t)], \quad (9)$$

where $\rho[R_t^e, \ln(M_t)]$ is the correlation of the excess return and the log SDF, and $\sigma(R_t^e)$ and $\sigma[\ln(M_t)]$ are the standard deviations of the excess return and log SDF respectively.

The log SDF in (6) is linear in consumption growth so the assumption of normality boils down to assuming that R_t^e and Δc_t have a bivariate normal distribution. We can then write (8)–(9) as

$$E(R_t^e) = \text{Cov}(R_t^e, \Delta c_t)\gamma, \text{ and} \quad (10)$$

$$E(R_t^e)/\sigma(R_t^e) = \rho(R_t^e, \Delta c_t) \times \sigma(\Delta c_t)\gamma, \quad (11)$$

This expression is convenient since it allows us to split the Sharpe ratio into three separate factors: the correlation of returns and consumption growth, the volatility of consumption growth, and the risk aversion. In particular, the first two factors can be estimated directly from data without knowing any preference parameters and the volatility is a constant in a cross-section (something we will use later).

Equation (11) shows that to get a positive expected excess return, the asset must be “risky” in the sense that it is positively correlated with consumption growth (negatively

---

⁷This does not imply (but is implied by) normally distributed gross returns, which would eventually violate limited liability. Even so, that would only happen once every 3000 years or so with the means and volatilities in Table 1. I just wonder if Psusennes I (Pharao 1044BC to 994BC) allowed limited liability. In any case, his gold mask (Egyptian Museum, Cairo) is second only to Tutanchamon’s.

⁸Stein’s lemma says that if x and y have a bivariate normal distribution and $h(y)$ is a differentiable function such that $E[|h'(y)|] < \infty$, then $\text{Cov}[x, h(y)] = \text{Cov}(x, y) E[h'(y)]$. This holds also for conditional distributions. In (7), $x = R_t^e$, $y = \ln(M_t)$, and $h(y) = \exp(y)$ so $h'(y) = M_t$.

correlated with the SDF). In that case, the asset tends to give a high return when marginal utility is low, so investors demand a risk premium. This effect is strong if consumption is very volatile and/or investors are strongly risk averse.

Table 1 shows that the stock market has a Sharpe ratio around 0.35–0.5 (annualized) and that the consumption growth rate has a standard deviation of around $1\% = 0.01$ (annualized). The lowest value of γ which makes equation (11) hold is then $\gamma = 35$, and this would only work if the stock market were a worthless hedge against consumption fluctuations so the correlation is one.⁹

The basic problem with the consumption-based asset pricing model is that investors enjoy a fairly stable consumption series (either because income is smooth or because it is easy/inexpensive to smooth consumption by changing savings), so only an extreme risk aversion can motivate why investors require such a high equity premium. This is the *equity premium puzzle* stressed by Mehra and Prescott (1985) (although they approach the issue from another angle—more about that later on).

Still higher values of γ are needed if the correlation in (11) is lower than one—as it certainly is. Table 1 shows that the correlation is around 0.1–0.3, where the higher value is for the correlation with Δc_{t+1} . If we want to be nice to the consumption-based model, then we can use the higher value. With a Sharpe ratio of 0.35 (the lowest value, once again to be nice) and a standard deviation of consumption of 0.01, (11) becomes $0.35 = 0.3 \times 0.01\gamma$, which requires $\gamma = 117$.

The low correlation therefore adds a *correlation puzzle* (see, for instance, Cochrane (2001)) to the equity premium puzzle: the consumption-based model makes the aggregate US stock market look pretty safe since almost all risk is non-systematic; it would indeed take a high risk aversion coefficient to motivate the historical risk premium.

If the consumption-based model cannot (with reasonable parameter values) explain the risk premium on the aggregate US stock market—what about other “aggregate” assets like US government bonds? Table 1 shows that the Sharpe ratio for bond returns is 0.12 for the long sample but as high as 0.27 for the shorter (and later) sample—but the sampling variability is so large so neither of these figures is significantly (on the 5% level) different

---

⁹This is related to the approach in Hansen and Jagannathan (1991). Rewrite (4) as $E(R_t^e)/\sigma(R_t^e) = -\rho(R_t^e, M_t)\sigma(M_t)/E(M_t)$, which shows that the following “volatility bound” must hold: $\sigma(M_t)/E(M_t) \geq |E(R_t^e)|/\sigma(R_t^e)$. Their approach is essentially to search in a set of returns for the portfolio with the largest Sharpe ratio and then study if a given model of the SDF satisfies the volatility bound. In the case considered here there is only one asset.

Figure 1: The equity premium puzzle for different samples. Subfigure a shows recursive estimations (longer and longer sample) of the Sharpe ratio of the aggregate US stock market and of the correlation with Δc_{t+1} times the standard deviation of Δc_{t+1} . Subfigures b and c show the same statistics calculated on a moving data window of 10 years and 5 years respectively. See Appendix A for details on data sources and transformations.

from zero. In any case, if we use the point estimates we need a risk aversion coefficient (γ) of around 10 or 25 to make equation (11) hold—which looks much more promising than for equity. However, the correlation of the bond return and consumption growth is very low, so the equity premium puzzle is also a bond premium puzzle.

2.2.1 The Equity Premium Puzzle over Time

The results for the two samples in Table 1 are somewhat different, which immediately raises the question of how sensitive the conclusions are to the sample period. *Figure 1* provides some answers by plotting the key statistics for different samples.

Figure 1.a shows recursive estimates of the Sharpe ratio of the aggregate US stock market and the Sharpe ratio predicted by the CRRA model when the risk aversion is

100, that is, $\rho(R^e, \Delta c) \times \sigma(\Delta c) \times 100$ from (11). The correlations are for Δc_{t+1} . The recursive estimation means that the results for 1957Q2 use data for 1947Q2–1957Q2, the results for 1957Q3 add one data point, etc. The results for 2001Q4 are therefore the same as those in the top panel of Table 1. Figure 1.b shows the same statistics, but estimated on a moving data window of 10 years. For instance, the results for 1980Q2 are for the sample 1971Q3–1980Q2. Finally, Figure 1.c uses a moving data window of 5 years.

Together these figures give the impression that there are fairly long swings in the data. This fundamental uncertainty should serve as a warning against focusing on the fine details of the data. It could also be used as an argument for using longer data series—provided we are willing to assume that the economy has not undergone important regime changes.

There are other interesting results in these figures. In particular, the actual Sharpe ratio and the prediction of the consumption-based model move in different directions—which is something we will return to later in discussing time-variation of risk premia (Section 2.5).

2.2.2 The Mehra-Prescott Approach

The Mehra and Prescott (1985) approach to analyzing the consumption-based asset pricing model is somewhat less direct than what we did above. However, it has had a large impact on the research community, so it is worth summarising.

The starting point is a Lucas (1978) economy where (aggregate) consumption equals the exogenous (aggregate) endowment. To see how this works, write the gross return as $R_t = (P_t + D_t)/P_{t-1}$, where P_t is the asset price and D_t the dividend (per asset) in period t . We can then write (1) as $P_{t-1} = E_{t-1}[M_t(P_t + D_t)]$. Iterating forward, applying the law of iterated expectations, ruling out bubbles, and using the CRRA SDF in (5) give

$$P_{t-1} = E_{t-1}[\beta(C_t/C_{t-1})^{-\gamma} D_t + \beta^2(C_{t+1}/C_{t-1})^{-\gamma} D_{t+1} + \dots]. \quad (12)$$

The price of a real one-period bond is found by setting $D_t = 1$ and all future dividends to zero. The price of a claim to the stream of future non-storable endowments—which is interpreted as a proxy of a broad stock market index—is found by setting D_t equal to C_t . With a time series process for consumption, and values of β and γ , it is straightforward to calculate asset returns.

Mehra and Prescott (1985) construct a time series process by postulating that consumption growth follows a two-state Markov process: C_t/C_{t-1} is either low or high and the probability of which depends on whether C_{t-1}/C_{t-2} was low or high. This makes it very simple to calculate the expectations in (12): there is only one state variable and it can only take two values. In spite of this, the process allows both autocorrelation and time-varying volatility. The model parameters are calibrated to fit the mean, standard deviation, and autocorrelation in consumption growth. With this machinery, Mehra and Prescott (1985) find that a very high value of γ is needed to explain the historical equity premium—basically because consumption is fairly smooth.

This approach is interesting, but it is a fairly indirect method for studying the equity premium, and it adds a few very strong assumptions, in particular, about the consumption process and that the stockmarket corresponds to a claim on future consumption. In any case, several extensions have been made, for instance to investigate if more extreme consumption growth processes (fat tails or “crash states”) can rescue the model—but so far with limited success (see, for instance, Salyer (1998) and Bidarkota and McCulloch (2000)).

2.3 The Riskfree Rate Puzzle

The CRRA utility function has the special feature that the intertemporal elasticity of substitution is the inverse of the risk aversion, that is, $1/\gamma$. The choice of the risk aversion parameter, for instance, to fit the equity premium, will therefore have direct effects on the riskfree rate.

A key feature of any consumption-based asset pricing model, or any consumption/saving model for that matter, is that the riskfree rate governs the time slope of the consumption profile. From the asset pricing equation for a riskfree asset (3) we have $E_{t-1}(R_{ft}) E_{t-1}(M_t) = 1$, which in the CRRA model becomes

$$E_{t-1}(R_{ft}) E_{t-1}[\beta(C_t/C_{t-1})^{-\gamma}] = 1. \quad (13)$$

Note that we must use the conditional asset pricing equation—at least as long as we believe that the riskfree asset is a random variable. A riskfree asset is defined by having a zero conditional covariance with the SDF, which means that it is regarded as riskfree at the time of investment ($t - 1$). In practice, this means a real interest rate (perhaps

approximated by the real return on a T-bill since the innovations in inflation are small), which may well have have a nonzero unconditional covariance with the SDF.¹⁰ Indeed, in Table 1 the real return on a T-bill is as correlated with consumption growth as the aggregate US stockmarket.

To make progress, assume that the log SDF and the log riskfree rate $r_{ft} = \ln(R_{ft})$ have a bivariate normal distribution—which is somewhat different from our earlier assumption, but should (once again) be seen as just a an approximation. We can then write (13) as¹¹

$$E_{t-1}(r_{ft}) + \sigma_{t-1}^2(r_{ft})/2 = -\ln(\beta) + \gamma E_{t-1}(\Delta c_t) - \gamma^2 \sigma_{t-1}^2(\Delta c_t)/2. \quad (14)$$

To relate this equation to historical data, we take unconditional expectations to get

$$E(r_{ft}) + E \sigma_{t-1}^2(r_{ft})/2 = -\ln(\beta) + \gamma E(\Delta c_t) - \gamma^2 E \sigma_{t-1}^2(\Delta c_t)/2. \quad (15)$$

Before we try to compare (15) with data, several things should be noted. First, the log gross rate is very close to a traditional net rate ($\ln(1+z) \approx z$ for small z), so it makes sense to compare with the data in Table 1. Second, we can safely disregard the variance terms since they are very small, at least as long as we are considering reasonable values of γ . Although the average conditional variances are not directly observable, we know that they must be smaller than the unconditional variances¹², which are very small in Table 1. In fact, the variances are around 0.0001 whereas the means are around 0.01–0.02.

According to (15) there are two ways to reconcile a positive consumption growth rate with a low real interest rate (2%/year and 1%–1.5% respectively in Table 1): investors may prefer to consume later rather than sooner ($\beta > 1$) or they are willing to substitute intertemporally without too much compensation ($1/\gamma$ is high, that is, γ is low). However, fitting the equity premium requires a high value of γ , so investors must be implausibly patient if (15) is to hold. For instance, with $\gamma = 25$ (which is a very conservative guess of what we need to fit the equity premium) equation (15) says $0.01 = -\ln(\beta) + 25 \times 0.02$ (ignoring the variance terms) for the long sample, which requires $\beta \approx 1.6$. The shorter

¹⁰As a very simple example, let $x_t = z_{t-1} + \varepsilon_t$ and $y_t = z_{t-1} + u_t$ where ε_t are u_t uncorrelated with each other and with z_{t-1} . If z_{t-1} is observable in $t-1$, then $\text{Cov}_{t-1}(x_t, y_t) = 0$, but $\text{Cov}(x_t, y_t) = \sigma^2(z_{t-1})$.

¹¹If $x \sim N(\mu, \sigma^2)$ and $y = \exp(x)$ then $E(y^k) = \exp(k\mu + k^2\sigma^2/2)$.

¹²Let $E(y|x)$ and $\text{Var}(y|x)$ be the expectation and variance of y conditional on x . The unconditional variance is then $\text{Var}(y) = \text{Var}[E(y|x)] + E[\text{Var}(y|x)]$.

sample gives a very similar number. This is the *riskfree rate puzzle* stressed by Weil (1989). The basic intuition for this result is that it is hard to reconcile a steep slope of the consumption profile and a low compensation for postponing consumption if people are insensitive to intertemporal prices—unless they are extremely patient (actually, unless they prefer to consume later rather than sooner).

Another implication of a high risk aversion is that the real interest rate should be very volatile, which it is not. According to Table 1 the standard deviation of the real interest rate is 1.5 times the standard deviation of consumption growth. From (14) the volatility of the (expected) riskfree rate should be

$$\sigma[E_{t-1}(r_{ft})] = \gamma \sigma[E_{t-1}(\Delta c_t)], \quad (16)$$

if the variance terms are constant. The standard deviation of expected real interest rate is γ times the standard deviation of expected consumption growth. We cannot observe the conditional expectations directly, and therefore not estimate their volatility. However, a simple example is enough to demonstrate that high values of γ are likely to imply counterfactually high volatility of the real interest rate.

Suppose that consumption growth is an AR(1) process. In that case, $\sigma[E_{t-1}(\Delta c_t)]$ equals $\rho(\Delta c_t, \Delta c_{t+1})\sigma(\Delta c_t)$, which from Table 1 is 0.22% for the long sample and 0.37% for the short sample.¹³ With $\gamma = 25$, (16) then implies that the standard deviation of the expected real interest rate should be around 5%–9%. Table 1 shows that the unconditional standard deviation, which is an upper bound (see footnote 12), is around 1.5%. This shows that an intertemporal elasticity of substitution of 1/25 is not compatible with the relatively stable real return on T-bills.

2.4 The Cross-Section of Returns

The previous section demonstrated that the consumption-based model has a hard time explaining the risk premium on a broad equity portfolio—essentially because consumption growth is too smooth to make stocks look particularly risky. However, the model *does* predict a positive equity premium, even if it is not large enough. This suggests that the model may be able to explain the relative risk premia across assets, even if the scale is wrong. In that case, the model would still be useful for some issues. This section takes a

---

¹³If $x_t = \alpha x_{t-1} + \varepsilon_t$, where ε_t is iid, then $E_{t-1}(x_t) = \alpha x_{t-1}$, so $\sigma(E_{t-1} x_t) = \alpha \sigma(x_{t-1})$.

closer look at that possibility by focusing on the relation between the Sharpe ratio and the correlation with consumption growth in a cross-section of asset returns.

The key equation is (11), which I repeat here for ease of reading

$$E(R_t^e)/\sigma(R_t^e) = \rho(R_t^e, \Delta c_t) \times \sigma(\Delta c_t)\gamma. \quad (11)$$

The volatility of the log SDF, $\sigma(\Delta c_t)\gamma$, is the same for all assets, so the cross-sectional variation in Sharpe ratios is due to cross-sectional variation in the correlation of the portfolio excess return and consumption growth, $\rho(R_t^e, \Delta c_t)$. We can therefore study the cross-sectional implication of (11) by estimating the linear regression equation

$$E(R_t^e)/\sigma(R_t^e) = a_0 + a_1\rho(R_t^e, \Delta c_t) + u, \quad (17)$$

where the observations are different portfolios. According to (11) the slope coefficient a_1 should be positive (since $\sigma(\Delta c_t)\gamma$ is) and the coefficient of determination R^2 should be high (unity, except for sampling variability).

In a sense, this regression is similar to the traditional cross-sectional regressions of returns on factors with unknown factor risk premia (see, for instance, Cochrane (2001) chap 12 or Campbell, Lo, and MacKinlay (1997) chap 6).

Figure 2 studies different cross-sections of assets by plotting the Sharpe ratio against the correlation with consumption growth. To save some space, only results for Δc_{t+1} are shown. To be honest, the corresponding results for Δc_t look even less flattering. *Figure 2.a* shows evidence for 10 US industry portfolios (telecom, utilities, etc) for the sample 1970Q1–2001Q4; see Appendix A for details on data. There seems to be a weak positive relation between the Sharpe ratio and the correlation with consumption. This is confirmed by the numbers at the bottom of the figure: the slope coefficient in the estimated regression equation (17) is 0.3, but it is not significantly different from zero at the 5% level (would have been marked by a star * as in *Figure 2.d*) and the R^2 is only 0.08. The result is no better if the odd portfolio with zero correlation (oil, gas, and coal) is excluded.

Another way to see essentially the same result is to compare the industry portfolios with the aggregate US market (marked by a large circle in the figure). In this figure, the market happens to have the highest correlation with consumption growth and should according to (11) therefore have the highest Sharpe ratio. In fact, six the ten portfolios have higher Sharpe ratios, of which three are considerably higher.

Figure 2: Sharpe ratios and consumption correlation, 1970–2001. Subfigure a shows Sharpe ratios and correlations with Δc_{t+1} for 10 US industry portfolios for 1970Q1–2001Q4. The large circle is the aggregate US stock market. The numbers at the bottom are slope coefficients and R^2 from a regression of the Sharpe ratio on a constant and the correlation. Significance at the 5% level is marked by a star *. Subfigures b–d are similar to subfigure a but for US dividend/price, size, and book-to-market portfolios respectively. Subfigures e–f are for international equity and bond portfolios respectively. See Appendix A for details on data sources and transformations.

There is also very little support for the model on US dividend/price and size portfolios (deciles of D/P and size respectively): the R^2 s are virtually zero in Figures 2.b–c. The model seems to work somewhat better on US book-to-market portfolios (Figure 2.d), but fails again on the international equity and bond portfolios (Figures 2.e–f).

It is clear that the consumption-based model has problems with explaining the cross-section of Sharpe ratios. The main problem is that there are considerable differences in Sharpe ratios, but not in correlations with consumption growth. To put it bluntly: systematic risk seems to be something else than covariance with consumption growth. However, the model is not a complete failure: 4 of 6 figures at least get a positive slope.

Figure 3: Sharpe ratios and market correlation, 1970–2001. These figures show Sharpe ratios plotted against the correlation with the aggregate US market excess return. See Figure 2 for details.

To get a perspective on the magnitude of the problems, I will compare with the traditional Capital Asset Pricing Model (CAPM, Lintner (1965), Mossin (1968), and Sharpe (1964)).

The CAPM is most easily derived by postulating that the SDF is an affine function of the market excess return, R_{mt}^e ,

$$M_t = a + bR_{mt}^e. \quad (18)$$

(With a constant riskfree rate, this is the same as assuming an affine function of the market return.). Apply the unconditional version of (4) on the market excess return and some other excess return, R_t^e , and combine the results to get the standard CAPM expression

$$E(R_t^e) = \frac{\text{Cov}(R_t^e, R_{mt}^e)}{\sigma^2(R_{mt}^e)} E(R_{mt}^e). \quad (19)$$

This says that the excess return on an asset is a regression coefficient (R_t^e regressed on

R_{mt}^e) times the market excess return.

For a direct comparison with the consumption-based model in (11), it is more convenient to write (19) as

$$E(R_t^e)/\sigma(R_t^e) = \rho(R_t^e, R_{mt}^e) \times E(R_{mt}^e)/\sigma(R_{mt}^e), \quad (20)$$

which says that the Sharpe ratio is a positive linear function (with a slope equal to the Sharpe ratio of the market, that is, around 0.5) of the correlation of the asset with the market.

Figure 3 illustrates (20). The figure has the same structure as *Figure 2*, except that the correlation is now with the excess return on the aggregate US market. These correlations are typically much higher than the correlations with consumption growth. However, it is hard to argue that the correlations with the stock market are much better at explaining the cross-section of Sharpe ratios: the slope coefficient is now positive in only 2 of 6 cases. In particular, CAPM gives odd results (negative coefficients) for the dividend-price and the book-to-market portfolios (see Fama and French (1993) for a detailed analysis of these portfolios).

The conclusion is that the consumption-based model is not good at explaining the cross-section of Sharpe ratios, but it is no worse than CAPM—if it is any comfort.

2.5 Time-Variation in Risk Premia

In contrast to the traditional interpretation of “efficient markets,” it has been found that excess returns might be somewhat predictable—at least in the long run (a couple of years). In particular, Fama and French (1988a) and Fama and French (1988b) have argued that future long-run returns can be predicted by the current dividend-price ratio and/or current returns.

Table 2 illustrates this by showing the R^2 from the regressions

$$R_{t+k}^e(k) = a_0 + a_1 x_t + u_{t+k}, \text{ where } x_t = D_t/P_t \text{ or } R_t^e(k), \quad (21)$$

where $R_t^e(k)$ is the annualized k -quarter excess return of the aggregate US stock market and D_t/P_t is the dividend-price ratio (see Appendix A for details on data).

It seems as if the dividend-price ratio has some explanatory power for future returns—at least for long horizons in the long sample (first panel, 1947–2001). For instance, the

Horizon in quarters, (k)	1	4	8	12	16	20
<u>1947–2001</u>						
D_t/P_t	0.02*	0.10*	0.13*	0.16*	0.20*	0.27*
$R_t^e(k)$	0.00	0.03	0.03*	0.04	0.13*	0.07*
<u>1970–2001</u>						
D_t/P_t	0.01	0.04	0.01	0.00	0.00	0.00
$R_t^e(k)$	0.00	0.06*	0.05	0.00	0.03	0.03

Table 2: Predicting long-run excess returns on the aggregate US stockmarket. This table shows results from regressing future k -quarter returns, $R_{t+k}^e(k)$, on the current dividend-price ratio (D_t/P_t) or the current k -quarter return, $R_t^e(k)$. The samples are 1947Q2–2001Q4 and 1970Q1–2001Q4. A star (*) denotes a significance at the 5% level, based on GMM/delta method inference and a Newey-West estimator with $k - 1$ lags. See Appendix A for details on data sources and transformations.

R^2 is around 0.20–0.25 for the 4–5 year horizons. However, the effect vanishes in the late sample (second panel, 1970–2001). The lagged return is a fairly weak predictor in both samples.

This evidence suggests that excess returns may perhaps have a predictable component, that is, that (ex ante) risk premia are changing over time. To see how that fits with the consumption-based model, assume that consumption growth is conditionally normally distributed. We then get a conditional Sharpe ratio as (compare with (11))

$$E_{t-1}(R_t^e)/\sigma_{t-1}(R_t^e) = \rho_{t-1}(R_t^e, \Delta c_t) \times \sigma_{t-1}(\Delta c_t)\gamma. \quad (22)$$

By multiplying both sides by $\sigma_{t-1}(R_t^e)$ we see that the conditional expected excess return should equal the conditional covariance times the risk aversion.

It is clear from the earlier Figures 1.a–c that the consumption-based model probably cannot generate plausible movements in risk premia. In those figures the conditional moments in (22) are approximated by estimates on different data windows (that is, different subsamples). Although this is a crude approximation, the results are revealing: the actual Sharpe ratio and the prediction of the consumption-based model move in different directions on all frequencies. In Figure 1.a (longer and longer sample) the actual Sharpe ratio is decreasing over time, whereas the prediction is mostly increasing. In Figure 1.b the actual Sharpe ratio (on 10-year subsamples) reaches a minimum during the late 1970s and

early 1980s—at the same time as the prediction is at maximum. The 5-year subsamples in Figure 1.c seems to pick up business cycle movements, and the two series clearly move in different directions.

2.6 Refinements of the Consumption-Based Model

There are many suggestions for how the problems with the consumption-based can be solved. One major strand proposes changes to the utility function; another to how we measure the volatility of returns and consumption.

If we focus on the utility function, we need a high risk aversion (to get a high equity premium), a high elasticity of intertemporal elasticity of substitution (to get a low and stable real interest rate), and maybe also time-variation of the risk aversion (to fit the time-variation in expected returns). This could be produced by using a functional form that explicitly separates risk aversion from the intertemporal elasticity of substitution as in Epstein and Zin (1989a) or by certain types of habit persistence models as in Campbell and Cochrane (1999).

If we instead focus on the measurement of volatility, we need a high consumption volatility (to get a high equity premium) and time-variation in the volatility of returns and/or consumption (to get time-variation in expected returns). The first could be achieved by taking into account the effects of uninsurable idiosyncratic shocks as in Mankiw (1986); the second by a model of time-variation of second moments (typically a GARCH or regime switching model) as in, for instance, Bansal and Lundblad (2000).

There are of course many unresolved issues. First, some of the proposed “fixes” have serious side effects. For instance, some of the habit persistence models create implausible implications for consumption smoothing and optimal fiscal policy (see, Lettau and Uhlig (2000), Lettau and Uhlig (2002), and Ljungqvist and Uhlig (2000)).

Second, it is unclear if these modifications of the SDF can improve the consumption-based model’s ability to account for the cross-section of Sharpe ratios discussed in Section 2.4, since the cross-sectional variation often depends on the correlation with consumption growth only—just as in the CRRA model. The next few sections demonstrate this for some well-known models of time nonseparable utility (Epstein and Zin (1989a)), habit persistence (Campbell and Cochrane (1999)), and idiosyncratic shocks (Constantinides and Duffie (1996)).

2.6.1 Epstein-Zin Utility

The basic idea of the recursive utility function in Epstein and Zin (1989b) is to form a certainty equivalent of future utility as $Z_t = [E_t(U_{t+1}^{1-\gamma})]^{1/(1-\gamma)}$ where γ is the risk aversion—and then use a CES aggregator function to govern the intertemporal trade-off between current consumption and the certainty equivalent: $U_t = [(1 - \delta)C_t^{1-1/\psi} + \delta Z_t^{1-1/\psi}]^{1/(1-1/\psi)}$ where ψ is the elasticity of intertemporal substitution.

Epstein and Zin (1989b) show that if all wealth is marketable, so the budget restriction can be written $W_{t+1} = R_{mt+1}(W_t - C_t)$ where R_{mt+1} is the market return (in a broad sense), then the Euler equation for an asset with return R_t is

$$1 = E_{t-1}[\beta^\theta (C_t/C_{t-1})^{-\theta/\psi} R_{mt}^{\theta-1} R_t], \text{ where } \theta = (1 - \gamma)/(1 - 1/\psi). \quad (23)$$

This gives rise to rather complicated pricing expressions, but it is straightforward to show that the basic mechanisms are as in the CRRA model, except that the elasticity of intertemporal substitution is no longer forced to be the inverse of the risk aversion.

To illustrate that (23) has (approximately) the same implications for risky assets as the CRRA model, assume that consumption-wealth ratio is constant (for instance, because the market return is iid). This effectively closes down all interaction between the risk and intertemporal substitution (see Svensson (1989) and Campbell (1993)). It follows from the budget constraint that consumption growth and the market return are proportional so (23) can be written

$$E_{t-1}[(C_t/C_{t-1})^{-\gamma} R_t] = \text{constant}, \quad (24)$$

which is exactly the same as the CRRA model. The Epstein-Zin model therefore inherits the problem with explaining the cross-section of Sharpe ratios. This is also approximately true when the consumption-wealth ratio is not too variable. Of course, if there are large predictable movements in the market return, then it is no longer true.

To see that the intertemporal substitution is governed by the parameter ψ , close down all systematic risk by assuming that the market return is a riskfree rate (set $R_{mt} = R_t = R_{ft}$ in (23)) and that log consumption and the log real interest rate, r_{ft} , have a bivariate normal distribution. We then get

$$E_{t-1}(r_{ft}) + \theta \sigma_{t-1}^2(r_{ft})/2 = -\ln(\beta) + \frac{1}{\psi} E_{t-1}(\Delta c_t) - \frac{\theta}{\psi^2} \sigma_{t-1}^2(\Delta c_t)/2. \quad (25)$$

This is the same relation between the real interest rate and consumption growth as in the CRRA case (see (14))—except that the elasticity of intertemporal substitution is no longer forced to be the inverse of the risk aversion.

This analysis shows that the Epstein-Zin model could help us to solve the riskfree rate puzzle (since we have a separate parameter for the elasticity of intertemporal substitution), but it can probably not do much with the equity premium (unless we are willing to accept a very high risk aversion) or the cross-sectional variation in Sharpe ratios.

2.6.2 Habit Persistence

The utility function in the habit persistence model of Campbell and Cochrane (1999) has the same functional form as the CRRA model, but the argument is the difference between consumption and a habit level, $C_t - X_t$, instead of just consumption. Variation in consumption and habit can both contribute to variation in marginal utility, which affects asset prices.

Campbell and Cochrane (1999) parameterize the habit in terms of the “surplus ratio” $S_t = (C_t - X_t)/C_t$, which measures how much aggregate consumption exceeds the habit. Since this ratio is external to the investor, marginal utility becomes $(C_t - X_t)^{-\gamma} = (C_t S_t)^{-\gamma}$. We can therefore write the log SDF as

$$\ln(M_t) = \ln(\beta) - \gamma(\Delta s_t + \Delta c_t), \quad (26)$$

where s_t is the log surplus ratio. The process for s_t is assumed to be a non-linear AR(1)

$$s_t = \phi s_{t-1} + \lambda(s_{t-1}) \Delta c_t + \text{constant}, \quad (27)$$

where $\lambda(s_{t-1})$ is a decreasing function of s_{t-1} which controls the sensitivity to (aggregate) consumption growth. The function $\lambda(s_{t-1})$ is always non-negative, so the log surplus ratio reacts positively to consumption growth (the main effect of letting λ depend on s_{t-1} is to make the risk premium move with the business cycle). Because of (27), a move in Δc_t will typically be accompanied by a move of Δs_t in the same direction, which increases the volatility of the SDF (26). In other words, a move in consumption drives the habit X_t in the opposite direction so $C_t - X_t$ varies a great deal.

Consumption growth is assumed to be conditionally normally distributed, so the con-

ditional Sharpe ratio is (compare with (9))

$$E_{t-1}(R_t^e)/\sigma_{t-1}(R_t^e) = -\rho_{t-1}[R_t^e, \ln(M_t)] \times \sigma_{t-1}[\ln(M_t)]. \quad (28)$$

From (26) and (27) we note that the innovation in the log SDF is $-\gamma[\lambda(s_{t-1}) + 1]$ times the innovation in Δc_t . Since $\lambda(s_{t-1})$ is positive and known in $t - 1$ this is very similar to the CRRA model, except that the conditional volatility of the log SDF (the last term in (28)) is scaled up. This could help explain a higher risk premium. In a way, it is very similar to just increasing the risk aversion γ (although Campbell and Cochrane (1999) show that their model do not suffer from the some of the drawbacks of the CRRA model with a high γ).

However, (28) has the same implication for the cross-sectional variation of Sharpe ratios as the CRRA model. The reason is that the conditional correlation in (28) is the same as in the CRRA model since the innovation in the log SDF is proportional to the innovation in consumption growth (and the factor of proportionality cancels in the correlation).

2.6.3 Idiosyncratic Risk

A quite different attempt to solve the problems with the consumption-based asset pricing model is to argue that aggregating individual investors into a (fictitious) representative investor is a non-trivial task. Investors who face idiosyncratic risk that they cannot insure against (incomplete markets) may be reluctant to hold risky assets. Since macro data averages out idiosyncratic components, the volatility of aggregate consumption would then underestimate true consumption volatility of a representative investor.

However, it is surprisingly hard to make this story stick. The first problem is that only shocks that cannot be hedged qualify as idiosyncratic shocks: hedging means averaging out. One way of hedging consumption (if not income) is to have a savings account (see Telmer (1993)). Moreover, income can sometimes be partially hedged by trading in existing assets if they are correlated with income. In practice, this means that idiosyncratic shocks are probably a quite smaller than the income shocks estimated from micro data.

The second problem is that idiosyncratic shocks may still not matter unless their distribution is strongly tied to state of the economy. To develop an example, note that the basic asset pricing equation (1) must hold for every investor (after all, it is his first order condition), so $E_{t-1}(R_t M_{jt}) = 1$ where M_{jt} is the SDF of investor j . The conditional

version of the asset pricing equation with normally distributed consumption growth (10) for investor j is then

$$E_{t-1}(R_t^e) = \text{Cov}_{t-1}(R_t^e, \Delta c_{jt})\gamma, \quad (29)$$

where Δc_{jt} is consumption growth of investor j .

We let the innovation in the consumption investor j have two components: an aggregate innovation (common to all investors) and an idiosyncratic innovation. The idiosyncratic shock cannot be correlated with existing assets (discussed above), so it contributes nothing to the conditional covariance in (29): the risk premium will only depend on the covariance of the return with the aggregate innovation—as in the CRRA model. In this case the idiosyncratic shocks do not matter for the risk premium (although they will typically decrease the real interest rate as investors increase the precautionary saving).

This result is quite general. For instance, it holds for the Campbell and Cochrane (1999) model discussed before, and in any other model with a lognormal SDF where the idiosyncratic component affects the log SDF linearly (this follows directly from (8)).¹⁴

To make idiosyncratic shocks matter for risk premia, they will either have to be something less than completely idiosyncratic (for instance, affecting a group of investors who collectively have market power; see Den Haan (2001)) or their distribution must depend on the aggregate shock (see Mankiw (1986) and Constantinides and Duffie (1996)).

In the latter class of models, a higher risk premium requires that the volatility of the idiosyncratic shocks is stochastic and that the innovations to this volatility and of aggregate consumption are negatively correlated: bad times mean more risk. To illustrate this, consider the CRRA model where the log SDF for investor j is

$$\ln(M_{jt}) = \ln(\beta) - \gamma \Delta c_{jt}. \quad (30)$$

Let the innovation in consumption growth for investor j be

$$\Delta c_{jt} - E_{t-1}(\Delta c_{jt}) = \varepsilon_t + u_{jt}, \quad (31)$$

where ε_t is the aggregate component and u_{jt} the individual component which has a variance which depends on the realization of ε_t .

---

¹⁴Lettau (2002) shows that this property holds for any distribution of the consumption innovation if utility is CRRA and the the distribution of the idiosyncratic component is not affected by the aggregate shock.

With (30)–(31) the asset pricing equation for investor j says that

$$0 = E_{t-1}(R_t^e M_{jt}) \quad (32)$$

$$= E_{R,\varepsilon,u}\{R_t^e \exp[\ln(\beta) - \gamma E_{t-1}(\Delta c_t) - \gamma \varepsilon_t - \gamma u_{jt}]\}. \quad (33)$$

The operator $E_{R,\varepsilon,u}$ is used instead of E_{t-1} to indicate that the expectations involves integrating over the random variables R_t^e , ε_t , and u_{jt} . Suppose that the distribution of u_{jt} conditional on the aggregate innovation, ε_t , is normal with zero mean and variance $v_{t-1} + \lambda(\varepsilon_t)$ where v_{t-1} is known in $t - 1$ and $\lambda(\varepsilon_t)$ depends on the aggregate shock. Integrating over u_{jt} then gives

$$0 = E_{R,\varepsilon}\{R_t^e \exp[\ln(\beta) - \gamma E_{t-1}(\Delta c_t) + \gamma^2 v_{t-1} - \gamma \varepsilon_t + \gamma^2 \lambda(\varepsilon_t)/2]\}, \quad (34)$$

where $E_{R,\varepsilon}$ indicates that we now need to integrate over R_t^e and ε_t . (This first integration works since the asset return is independent of the realization of the idiosyncratic shock, even if not of its variance.)

I follow Lettau's (2002) approximation of assuming that ε_t and $\lambda(\varepsilon_t)$ have a bivariate normal distribution. This is only an approximation since it does not rule out negative variances, but it has the advantage of giving a very straightforward expression. The term in square brackets in (34) can then be interpreted as a normally distributed log SDF of the investor, which allows us to use the decomposition of the Sharpe ratio in (28) which I repeat here for easy of reading

$$E_{t-1}(R_t^e)/\sigma_{t-1}(R_t^e) = -\rho_{t-1}[R_t^e, \ln(M_t)] \times \sigma_{t-1}[\ln(M_t)]. \quad (28)$$

In this model, the conditional volatility and correlation are

$$\sigma_{t-1}[\ln M_t] = \gamma \sigma_{t-1}[\varepsilon_t - \gamma \lambda(\varepsilon_t)/2] \quad (35)$$

$$-\rho_{t-1}[R_t^e, \ln(M_t)] = \rho_{t-1}[R_t^e, \varepsilon_t - \gamma \lambda(\varepsilon_t)/2]. \quad (36)$$

Note that setting $\lambda(\varepsilon_t) = 0$ gives the same expressions as without idiosyncratic shocks. In fact, if $\lambda(\varepsilon_t)$ is a constant, that is, if the volatility of the idiosyncratic shocks do not depend on the aggregate shock, then we also get the same expressions as without idiosyncratic shocks.

If $\lambda(\varepsilon_t)$ is decreasing in ε_t so idiosyncratic risk is larger in bad times, then the volatil-

ity in (35) is larger than without idiosyncratic shocks. This could help explaining the equity premium puzzle—although there are doubts about whether the idiosyncratic risk is sufficiently volatile to be quantitatively important (see Lettau (2002) and Cogley (1998)).

However, this mechanism is unlikely to help the model to explain the cross-sectional variation in Sharpe ratios—which depends only on the correlation (36): as a first-order approximation this correlation is the same as when $\lambda(\varepsilon_t)$ is constant. For instance, if $\lambda(\varepsilon_t) = a + b\varepsilon_t$ with $b < 0$, then this is exactly so. This gives the same implications for the cross-sectional dispersion of Sharpe ratios as the model without idiosyncratic shocks—at least unless the idiosyncratic dispersion is a very non-linear function of the aggregate state.

An additional aspect of models with idiosyncratic risk is that they allow us to study the effects of transaction costs in a serious way: heterogeneous agents with idiosyncratic costs trade, representative agents do not. The first effect of transaction costs is to transform the first order condition of an investor to a set of inequalities $1/(1 + \tau) \leq E_{t-1}(R_t M_{jt}) \leq 1 + \tau$ where τ is the (proportional) cost of a round-trip (the upper limit is for buying the asset in $t - 1$ and selling in t ; the lower limit is for the opposite). This has been used to modify the Hansen and Jagannathan (1991) bounds (see He and Modest (1995) and Luttmer (1996)) with mixed results. One of the key assumptions in that analysis seems to be the assumption length of the period: a 0.5% transaction cost is quite substantive compared to a 1% monthly (expected) return, but small compared to a 12% annual return. We therefore need to find the equilibrium (including the endogenous trading frequency) in order to analyse the importance of transaction costs. So far, the findings indicate that the effect on prices is relatively small (unless borrowing/lending is very costly) but that the effect on turnover is large (see Heaton and Lucas (1996) and Vayanos (1998)).

3 Assets in Simulation Models

This part of the chapter discusses how different types of assets are incorporated into simulation models. The emphasis is on discussing the key properties of the assets when they are priced by a consumption-based model.

I first analyse a number of different asset classes: short- and long-lived claims on aggregate consumption, options and levered claims, as well as short and long real/nominal interest rates. To economise on space, I leave out some important asset classes like foreign

exchange and real estate (see, for instance, Smith and Wickens (2002), Kollman (2001), Zimmermann (1999) and Hendershott (1998)). At the end, I discuss some other important features of how assets are modelled: time-variation in risk premia, the effect of non-Walrasian labour contracts, and the effect of adding an asset to an existing model.

The analysis has the same basic approach as in the first part of the chapter. First, the model consists of the basic asset pricing equation (1) and the consumption-based stochastic discount factor (5). In practice, this means that we are studying equilibrium conditions on the joint behaviour of consumption growth and returns on assets that can be freely traded. This does not assume that markets are complete. However, adding an asset to a model with incomplete markets may change the consumption process, so the results we derive are only valid for assets that could be traded when the equilibrium was determined (see Section 3.8 for a more detailed discussion). Second, analytical pricing relations are derived by making assumptions about the distribution of returns and consumption and sometimes also by approximating complicated expressions.

Although this approach is useful for providing intuition for how different types of assets will work in a simulation model, it is clearly only a first step. In particular, this chapter never formulates and solves a full model. To do that, the other contributions to this book, the references found below, and the text by Altug and Labadie (1994) should all be very useful.

3.1 Preliminaries

Many simulation models are close to log-linear and are typically fed with normally distributed shocks. This means that most variables (including returns) are approximately lognormally distributed—sometimes exactly so. I therefore derive new expressions for risk premia and Sharpe ratios (the expressions in previous sections assumed normally, not lognormally, distributed returns). These expressions will be used repeatedly in the following sections.

The asset pricing equation with the CRRA SDF is $1 = E_{t-1}\{\beta \exp[-\gamma \Delta c_t + r_t]\}$, where $r_t = \ln(R_t)$ is the log return. If these variables have a bivariate normal distribution, it follows (see footnote 11) that the risk premium and Sharpe ratios are

$$E_{t-1}(\tilde{r}_t^e) = \text{Cov}_{t-1}(\gamma \Delta c_t, r_t^e), \text{ and} \quad (37)$$

$$E_{t-1}(\tilde{r}_t^e)/\sigma_{t-1}(r_t^e) = \rho_{t-1}(\Delta c_t, r_t^e) \times \sigma_{t-1}(\Delta c_t)\gamma. \quad (38)$$

The tilde is here to indicate that the risk premium is defined as $E_{t-1}(\tilde{r}_t^e) = E_{t-1}(r_t^e) + \sigma_{t-1}^2(r_t)/2$, that is, the expected excess log return over the real interest rate, adjusted for Jensen's inequality. This definition of an excess return is perhaps slightly non-standard, but it captures all the aspects of interest here.

In the following sections I will make use of the fact that only innovations matter for the risk premia, since it is only innovations that are risky. For the log excess return, this means that only the innovation in the risky asset payoff matters.¹⁵

3.2 Claims on Aggregate Consumption

This section analyses the pricing of claims on aggregate consumption. Many models have considered such assets—sometimes as an approximation of equity (see, for instance, Lucas (1978) and the references in Section 2.2.2). The main purpose of the discussion is to highlight that the risk premium on long-lived consumption claims depends crucially on the autocorrelation of consumption growth.

We first discuss a one-period claim and then move on to multiperiod claims. The real price in $t - 1$ of a claim on aggregate consumption in t satisfies the following asset pricing equation for CRRA utility (see (12))

$$P_{ct-1} = E_{t-1}[\beta(C_t/C_{t-1})^{-\gamma} C_t] \text{ or} \quad (39)$$

$$\ln(P_{ct-1}) = \ln(C_{t-1}) + \ln(\beta) + (1 - \gamma) E_{t-1}(\Delta c_t) + (1 - \gamma)^2 \sigma_{t-1}^2(\Delta c_t)/2, \quad (40)$$

where the second line exploits lognormality. The payoff in t is obviously C_t , so the return is C_t/P_{ct-1} . Instead of evaluating this by brute force (it is not very hard, just messy), we use (37). Note that the innovation in the excess log return of the consumption claim, r_{ct}^e , equals the innovation in consumption growth so (37) gives

$$E_{t-1}(\tilde{r}_{ct}^e) = \sigma_{t-1}^2(\Delta c_t)\gamma. \quad (41)$$

This asset has the highest possible Sharpe ratio ($\sigma_{t-1}(\Delta c_t)\gamma$) since the return is perfectly negatively correlated with the SDF—this is yet another version of the Hansen and Jagannathan (1991) bound. This means that if the model cannot generate a high Sharpe ratio for the one-period consumption claim, then it cannot do so for any other asset either.

---

¹⁵The log excess return, r_t^e is the log payoff minus the lagged log asset price (measured in real terms) minus the log real interest rate. Only the first component has an innovation.

Figure 4: Time line of $t - 1$ claims on C_t and C_{t+1} .

Most simulation models have rather focused on long-lived claims on a consumption stream, that is, a portfolio of different claims. What is the expected return on holding such a portfolio between two periods (the holding period return)?

It is easy to get the basic intuition by considering a claim that gives aggregate consumption in t and $t + 1$. When traded in $t - 1$ this is a portfolio of two claims: a claim on C_t (discussed above) and a claim on C_{t+1} . I will call the latter a “ C_{t+1} strip.” The time lines of these assets are shown in Figure 4. The holding period return (between $t - 1$ and t) on this portfolio is clearly the sum of the holding period returns on its two components.

We already know the (holding period) return on the one-period claim from (41), so I now turn to the C_{t+1} strip. This strip becomes a one-period claim in t and its price will then be like in (40) but with time subscripts advanced one period

$$\ln(P_{ct}) = \ln(C_t) + \ln(\beta) + (1 - \gamma) E_t(\Delta c_{t+1}) + (1 - \gamma)^2 \sigma_t^2(\Delta c_{t+1})/2. \quad (42)$$

This is the only term in the excess log holding period return (between $t - 1$ and t) with an innovation. Applying (37) therefore gives the expected holding period return (in excess of a real interest rate) of the strip

$$E_{t-1}(\tilde{r}_{ht}^e) = \gamma \sigma_{t-1}^2(\Delta c_t) + \gamma(1 - \gamma) \text{Cov}_{t-1}[\Delta c_t, E_t(\Delta c_{t+1})], \quad (43)$$

provided Jensen’s inequality term is constant (or at least uncorrelated with the innovation in consumption growth).

The second term in (43) is zero if consumption growth is not autocorrelated¹⁶ or if the

---

¹⁶ $\text{Cov}_{t-1}[\Delta c_t, E_t(\Delta c_{t+1})] = \text{Cov}_{t-1}(\Delta c_t, \Delta c_{t+1})$ since $\Delta c_{t+1} - E_t(\Delta c_{t+1})$ is not correlated with anything known in t .

utility function is logarithmic $\gamma = 1$. In this case, the strip has the same risk premium as the one-period claim in (41). This should come as no surprise since lack of autocorrelation makes time irrelevant and log utility makes the discount rate effect perfectly balance the “dividend” effect of any news.¹⁷

With positive autocorrelation in consumption (as Table 1 suggests), the strip has a lower expected return than the one-period claim—provided the risk aversion is larger than unity. The intuition is as follows. The price in (42) is high when expected consumption growth, $E_t \Delta c_{t+1}$, is low: although the expected payoff is low, it will be worth a lot since it comes at a time of scarcity (with log utility these effects cancel out). Now, if $E_t \Delta c_{t+1}$ is low, then period t consumption growth, Δc_t , is typically also low (they are positively correlated). This means that the strip generates a high holding period return in t (a capital gain) when consumption is scarce (high marginal utility): it is a fairly good hedge and will therefore not require a high risk premium. It is even conceivable that it could have a negative risk premium.

This extends to longer-lived claims on streams of consumption. Such assets will have a fairly low (high) risk premium if consumption is positively (negatively) autocorrelated for most leads/lags (see Campbell (1986)).

3.3 Options and Levered Claims

This section discusses options and option-like assets. Fairly few simulation models have studied traditional options directly, but some models have dealt with option-type assets like levered equity. However, simulation models could be a useful framework for studying and solving option pricing problems. I will illustrate this by considering a very simple case: a European call option that expires next period.

On the expiry date, the owner of a European call option has the right (not the obligation) to buy a prespecified asset for a prespecified (strike) price, X . If the asset payoff happens to be Y_t , then the payoff of the option is zero if $Y_t < X$ (the option is not exercised) and $Y_t - X$ otherwise.

With the CRRA SDF, the price of the call option must be

$$\text{Call option price}_{t-1} = E_{t-1}[\beta(C_t/C_{t-1})^{-\gamma} \max(0, Y_t - X)]. \quad (44)$$

---

¹⁷Abel (1999) assumes iid consumption growth, but reintroduces autocorrelation in the SDF through habit persistence.

If consumption growth and the log asset payoff, $\ln(Y_t)$, have a bivariate normal distribution, then the solution to this equation is the Black and Scholes (1973)/Merton (1973) formula (see Huang and Litzenberger (1988) or Söderlind and Svensson (1997) for straightforward and not at all tedious calculations). This result is quite natural when we compare with the dynamic hedging approach typically used for deriving the Black-Scholes formula, since both approaches fundamentally rely on lognormality of the asset payoff.¹⁸ This framework has been used by, for instance, Söderlind (2001a) to study the implications of monetary policy on bond options in a simple analytical real business cycle model.

To get intuition for the results on option and option-like assets, it is convenient to approximate the nonlinear payoff $\max(0, Y_t - X)$ with a power function

$$Z_t = Ay_t^\lambda, \text{ with } \lambda = 1/[1 - X/E_{t-1}(Y_t)] \text{ for } X < E_{t-1}(Y_t). \quad (45)$$

Abel (1999) shows that this approximation gives a coefficient of variation (standard deviation divided by the mean) of the payoff which is very close to the option payoff—at least for reasonably low values of X (so the probability of exercise is high). This approximate payoff is particularly convenient when Y_t is lognormally distributed since the lognormality carries over to Ay_t^λ : if $\ln(Y_t)$ is normally distributed, so is $\ln(A) + \lambda \ln(Y_t)$.

Figure 5 illustrates how (45) works for a payoff with the same distribution as US annual consumption growth (mean of 2% and a standard deviation of 1%), where the exercise price, X , is such that there is a 93% chance of exercise. The distribution of Y_t is also shown (it is lognormal, but looks almost like a normal distribution since the mean is so far from zero). The constant A in the power function payoff is chosen so that both payoffs have the same means of around 1.5%. It is straightforward to calculate (using the facts in footnote 18) that both the option payoff and the approximation have standard deviations close to 1%. We now use this approximation to study the implications of firm debt.

Many simulation models assume that the equity value equals the full value of a firm. However, most firms are financed by both equity and debt. Debt-to-assets (leverage) ratios vary a good deal, but values around 0.5 seems to be fairly typical for historical US data. Since equity is not a claim on the full value of the firm, only the residual claim, it has

---

¹⁸The result is easily derived by using the following facts: if $x \sim N(\mu, \sigma^2)$ and $y = \exp(x)$ then $E(y^k) = \exp(k\mu + k^2\sigma^2/2)$ and $E(y^k|y > a) = E(y^k)\Phi(r\sigma - a_0)/\Phi(-a_0)$, where $a_0 = [\ln(a) - \mu]/\sigma$; see Johnson, Kotz, and Balakrishnan (1994).

Figure 5: Example of option payoff and power approximation. This figure shows the payoff of a call option with strike price X and a power approximation. The value of the underlying asset is the random variable y with probability density function $\text{pdf}(y)$.

option-like features which is likely to make it more risky.

To illustrate the main effect of leverage, I compare the one-period claim on consumption discussed in Section 3.2 with a levered claim whose payoff is $\max(0, C_t - X)$ where X is the debt. Leverage reduces the payoff by subtracting a fixed amount (the debt), except that limited liability means that the payoff cannot be negative. A zero payoff (a return of -100%) is awfully bad for investors, so the levered claim will be considered very risky.

From (45) the levered claim has approximately the payoff AC_t^λ . This means that the innovation in the log payoff (which is what matters for the risk premium) is λ times the innovation in unlevered claim. It then follows directly from (41) that the risk premium of the levered claim, $E_{t-1}(\tilde{r}_{z_t}^e)$, is λ times the risk premium on the unlevered claim

$$E_{t-1}(\tilde{r}_{z_t}^e) = \lambda \sigma_{t-1}^2 (\Delta c_t) \gamma. \quad (46)$$

The levered claim has a higher risk premium since it is more volatile ($\lambda > 1$), but the Sharpe ratio is the same.

If we interpret $X/E_{t-1}(Y_t)$ as the leverage ratio which is (on average) around 0.5, then Abel's suggestion in (45) gives $\lambda = 2$. With this number the risk premium on levered claims should be twice as high as the risk premium on unlevered claims.

It is more difficult to model levered long-lived claims, since the claim can only go

“bankrupt” once, so numerical methods are called for. A few simulation models have incorporated leverage and arrived at somewhat different conclusions. For instance, Rouwenhorst (1995) finds a very small effect of leverage, while Jermann (1998) finds that it adds around one percentage point to the equity premium. The main reason for the difference seems to be that the firm value is very stable in Rouwenhorst’s model.

3.4 The Real Yield Curve

This section studies the yield curve of real interest rates. The price of a k -period real bond (with a unit payoff k periods ahead) must satisfy the asset pricing equation (see (12))

$$B_{kt-1} = E_{t-1}[\beta^k (C_{t+k-1}/C_{t-1})^{-\gamma}]. \quad (47)$$

I will consider two interest rates that can be calculated from this expression. First, the one-period (log) interest rate $r_{1t-1} = \ln(1/B_{1t-1})$, is the return on holding a real bond from $t - 1$ to t . Note that (according to convention) it carries the date of the investment (when the interest rate is agreed on), not the date of payoff as for most other assets. By exploiting the lognormality we get

$$r_{1t-1} = -\ln(\beta) + \gamma E_{t-1}(\Delta c_t) - \gamma^2 \sigma_{t-1}^2 (\Delta c_t)/2. \quad (48)$$

Real interest rates are time-varying only if expected consumption growth is (disregarding time-variation in Jensen’s inequality term).

Second, we can create a portfolio of one-period and two-period real (inflation-indexed) bonds to guarantee the investor a known (as of $t - 1$) interest rate on an investment between t and $t + 1$, that is, a real forward rate.

The forward-spot parity for an asset without intermediate dividends says that the contracted forward price (to be paid next period) is

$$\text{Forward price} = \frac{\text{Spot price}}{\text{Price of bond maturing next period}}.$$

The intuition is that the forward contract is like buying the asset today, but on credit.

The “spot way” of making sure that we get one unit (of the basket of goods) in $t + 1$ is to buy a two-period bond in $t - 1$ at the price B_{2t-1} . The price of a bond maturing next period is B_{1t-1} , so the forward price is B_{2t-1}/B_{1t-1} . This is illustrated in Figure 6.

Figure 6: Time line of how to get different interest rates.

We then define the forward rate as the rate of return on this investment: since the bond pays off one unit in $t + 1$ and the investment was B_{2t}/B_{1t} in t , the gross forward rate is B_{1t-1}/B_{2t-1} . From (47) we calculate the log forward rate, $f_{t-1} = \ln(B_{1t-1}/B_{2t-1})$, as

$$f_{t-1} = -\ln(\beta) + \gamma E_{t-1}(\Delta c_{t+1}) - \gamma^2 \sigma_{t-1}^2(\Delta c_{t+1})/2 - \gamma^2 \text{Cov}_{t-1}(\Delta c_t, \Delta c_{t+1}). \quad (49)$$

Both the short interest rate and the forward rate can vary over time. In a one-factor model of the yield curve (see, for instance, Vasicek (1977), Cox, Ingersoll, and Ross (1985), and Backus, Foresi, and Zin (1998)) they move in parallel. In (48)–(49) parallel movements require that $E_{t-1}(\Delta c_t)$ and $E_{t-1}(\Delta c_{t+1})$ are affine functions of a single variable, *the factor*. For instance, in the special case where expected consumption is an AR(1) with positive autocorrelation, then the real interest rates would actually follow a discrete version of Vasicek's (1977) model. That holds approximately in many simple real business cycle models, since their dynamics are well captured by a single dynamic state variable (see Cogley and Nason (1993)). This may sound awfully restrictive, but from the perspective of modelling the real yield curve, it may be acceptable. It still a matter of debate of whether the one-factor model can be rejected or not.

We now turn to studying the *average slope of the yield curve*, that is, the average over time. The risk premium is the forward rate (49) minus the expected (as of $t - 1$) future short rate (48) (but with time subscripts forwarded one period)

$$f_{t-1} - E_{t-1}(r_{1t}) = -\gamma^2 \text{Cov}_{t-1}(\Delta c_t, \Delta c_{t+1}) + \gamma^2 [\sigma_t^2(\Delta c_{t+1}) - \sigma_{t-1}^2(\Delta c_{t+1})]/2, \quad (50)$$

provided future conditional volatilities are known in advance.

The forward rate will be lower than the expected future short rate if consumption growth is positively autocorrelated (as Table 1 suggests), disregarding the Jensen's inequality terms. The intuition is the same as in Section 3.2: the forward contract will generate a capital gain in t if $E_t(\Delta c_{t+1})$ is low (the one-period interest rate is then low, so the price of a bond is high) and this will typically happen when consumption is scarce in t (positive autocorrelation).

This extends to longer interest rates, so the real yield curve will slope downwards (upwards) if consumption growth is positively (negatively) autocorrelated for most leads/lags. In other words, with positive autocorrelation in consumption growth, the long real (or inflation-indexed) bond is a good hedge and therefore “safer” than the strategy of rolling over short interest rates (see Campbell (1986), Backus, Gregory, and Zin (1989), Backus, Foresi, and Zin (1998), and Campbell and Viceira (2002)).

3.5 Money Supply and the Nominal Interest Rates

This section discusses the effect of monetary policy on nominal interest rates. Simulation models with nominal variables typically introduce money through a cash-in-advance constraint, by sticking real money balances into the utility function, or by a shopping-time technology (see, for instance, Labadie (1989) and Giovannini and Labadie (1991) for asset pricing analysis with cash-in-advance constraints). In many cases, the mere introduction of money has no effect on the real equilibrium, and in most other cases the effect is only temporary and negligible (see, for instance, Cooley and Hansen (1989)).

It is only when nominal rigidities are present that the real variables, including real assets, are affected—and this also gives monetary policy an important role for asset prices. For instance, a policy that aims at stabilizing output is likely to achieve some short-run stabilization of consumption, which may affect risk premia in general. On the other hand, exogenous monetary policy shifts add another source of uncertainty (as argued by Friedman and many others) which may work in the opposite direction. In any case, nominal bonds are likely to be more directly affected by monetary policy than other assets, so this section focuses on the nominal yield curve.

The real price of a nominal bond is as in (47) except that the real payoff is not unity but one divided by the nominal price index in $t + k - 1$, Q_{t+k-1} . Multiplying by today's

price level, Q_{t-1} , converts the real price to a nominal bond price

$$B_{kt-1}^{\$} = E_{t-1}[\beta^k (C_{t+k-1}/C_{t-1})^{-\gamma} Q_{t-1}/Q_{t+k-1}]. \quad (51)$$

I assume that log consumption and the inflation rate, $\pi_t = \ln(Q_t/Q_{t-1})$, has a bivariate normal distribution. The (log) one-period nominal interest rate, $i_{1t-1} = \ln(1/B_{1t-1}^{\$})$, is then

$$i_{1t-1} = -\ln(\beta) + E_{t-1}(\gamma \Delta c_t + \pi_t) - \sigma_{t-1}^2(-\gamma \Delta c_t - \pi_t)/2 \quad (52)$$

Using the real interest rate (48) gives

$$i_{1t-1} - E_{t-1}(\pi_t) - r_{1t-1} = -\text{Cov}_{t-1}(\gamma \Delta c_t, \pi_t) - \sigma_{t-1}^2(\pi_t)/2, \quad (53)$$

which is a modern version of the Fisher equation: the nominal interest rate equals the sum of expected inflation, the real interest rate and an inflation risk premium (and a Jensen's inequality term).

The Fisher equation (53) shows that the *movements in the nominal interest rate* come from at least two sources: expected inflation may change and the real interest rate too. Of course, in many traditional macro models the real interest rate is very stable so movements in the nominal interest rate reflect similar movements in inflation expectations (the “Fisher effect”).

The modelling of monetary policy is certainly crucial for the movements nominal interest rate. In models with optimal monetary policy it is often found that it is necessary to put interest rate volatility in the loss function of the policy maker (see, for instance, Söderlind (2001b), and Söderlind (1999)); otherwise the interest rate becomes much too volatile compared with data.¹⁹

In models with exogenous (non-optimal) monetary policy it can easily happen that the implied nominal interest rate has strange dynamics. As an example, consider what could happen in the model used in Cooley and Hansen (1989) and Cooley and Hansen (1995). The utility function is logarithmic ($\gamma = 1$), so the real interest rate equals the expected consumption growth plus a constant (see (48)). The cash-in-advance constraint makes the nominal value of consumption equal the exogenous money supply. Together, these facts mean that $E_{t-1}(\pi_t) + r_{1t}$ on the left hand side of (53) equals the expected money supply

---

¹⁹Several articles have also analysed the effects on real aggregates and the term structure of monetary policy shocks; see, for instance, Rudebusch (1998) and Cochrane and Piazzesi (2002b).

growth—and this carries over to the nominal interest rate. The money supply growth is modelled as an exogenous AR(1) process with positive autocorrelation, so this behaviour is inherited by the nominal interest rate (see footnote 13). It might seem innocent to change the autocorrelation coefficient to zero (so money supply becomes a random walk), but this would actually imply a constant nominal interest rate, which is slightly odd.

We now turn to discussing *the risk premium on the nominal bond* in (53). It is positive if the covariance of consumption growth and inflation is negative (which seems to be the case in US data). The reason is that the real return on the nominal bond (nominal return minus inflation) then tend to be low when consumption is scarce (see Cochrane and Piazzesi (2002a) for a recent empirical study of bond risk premia).

This risk premium is likely to depend on many modelling choices. With predetermined inflation (which is often used in New Keynesian models), there is no risk premium at all. Monetary policy is also likely to affect the risk premium. For instance, in new Keynesian models with some kind of Phillips curve (sticky, but forward looking, price setting, see, for instance, Clarida, Galí, and Gertler (1999)) output (consumption) and inflation move in the same direction if demand shocks (random movements in the time preference rate) dominate, but in opposite directions if supply shocks (random movements in productivity or factor input prices) dominate. If the monetary policy objective is to minimize some linear combination of variances of output and inflation without any regard to interest rate volatility, then it is well known that monetary policy should neutralize all demand shocks but only some of the supply shocks. This will tend to make risk premium positive.

3.6 Time-Variation in Risk Premia

This section briefly discusses how a simulation model can generate time-variation in risk premia. Recall from (22) that the Sharpe ratio in the CRRA model is

$$E_{t-1}(R_t^e) = \text{Cov}_{t-1}(R_t^e, \Delta c_t) \quad (54)$$

$$= \rho_{t-1}(R_t^e, \Delta c_t) \times \sigma_{t-1}(\Delta c_t) \times \sigma_{t-1}(R_t^e)^{-1}. \quad (55)$$

To get movements in risk premia in the consumption-based model, the conditional covariance of the excess return and consumption growth must be changing over time. Few simulation models generate this because of a combination of two features: the shocks are

typically homoskedastic and the dynamics is close to log-linear (so consumption growth and log returns are linear transformations of the shocks). The result is that all expected asset returns will (almost) move in parallel with the expected riskfree rate.

This can be changed by introducing heteroskedastic shocks. However, from (55) we see that time-variation in the volatility of returns only (with constant correlations) will make expected return move, but not Sharpe ratios. Similarly, time-variation in consumption volatility leads to proportional increases of the expected return and Sharpe ratios for all assets.

3.7 The Value of a Firm: Labour Contracts and Installation Costs

This section discusses how non-Walrasian labour contracts and installation costs (for new capital) affects the equity value. We put aside the issue of how the total value of a firm is split between equity holders and creditors in order to focus on the value of their combined stake. The traditional way of modelling a firm is to assume that it hires labour and finances investment through retained earnings (no debt). Any extra cash flow is distributed as dividends

$$D_t = Y_t - w_t L_t - I_t, \quad (56)$$

where Y_t is output, w_t the real wage rate, L_t labour input, and I_t investment costs. In reality most firms seem to smooth dividends, so D_t should probably not be directly compared with dividends, but with cash flows or earnings. Typically, the number of shares is kept constant and normalized to unity.

Most simulation models have a Cobb-Douglas production function $Y = Z_t K_t^\alpha L_t^{1-\alpha}$, where Z_t is the technology level, K_t capital stock, and L_t is labour input. The first order condition for hiring labour is that the real wage rate equals the marginal product of labour, $w_t = (1 - \alpha)Y/L_t$. This does not require a competitive labour market, only that the firm can hire freely at the going wage rate (“right to manage”). With this assumption, we can write (56) as

$$D_t = \alpha Y_t - I_t. \quad (57)$$

As side remark, we note that a high capital share, α , makes dividends fairly similar to consumption (the aggregate resource constraint is $Y_t = I_t + C_t$), so the consumption claim might be a good proxy for equity in that (admittedly unrealistic) case. In other cases, the dynamics of dividends and consumption can be quite different but many models still seem

to generate too stable dividends (see, for instance, Rouwenhorst (1995)). The basic reason is that neither aggregate output nor aggregate investments are nearly as volatile as cash flows and earnings on listed firms.

One possibility of making the cash flow process more realistic is to allow for labour hoarding (or other types of labour contracts that stabilizes the wage bill) and variation in the capacity utilization rate. Labour hoarding would make the wage bill relatively constant, rather than proportional to output as above. As a consequence, the cash flow would fluctuate more in response to movements in output. For instance, Dantihine and Donaldson (2002) study the implications different types of labour contracts in a setting where the priority status of wage claims creates an operational leverage (see also Section 3.3 for a discussion of leverage)—and find that this can create a sizeable risk premium.

It has also been suggested that investments should be subject to installation costs. Equations (56) and (57) are still true (I_t is the investment cost, not the gross addition of capital), but the behaviour of output, investment, and consumption will change. In particular, installation costs will make consumption more volatile since it makes consumption smoothing more costly (see, for instance, Jermann (1998) and Lettau (2000)). For instance, when it is prohibitively expensive to change the investment from some steady state level, then there is no consumption smoothing. Also, cash flows in (57) will react stronger to output than before, since investments will no longer move in the same direction as output. Both these effects tend to increase the equity premium, although it should be admitted that increased consumption volatility is something of a cheat since that is not a free parameter. (It could also be noted that installation costs leads to movements in Tobin's q , which seems to be a salient feature of data.)

3.8 Adding Assets to a Model

This section discusses if we really need fully-fledged equilibrium models to analyse the pricing of assets. The asset price in (12) is $P_t = E_t[\sum_{s=1}^{\infty} \beta^s (C_{t+s}/C_t)^{-\gamma} D_{t+s}]$, where D_{t+s} is the “dividend” in $t + s$. On a mechanical level, we probably need some kind of model to calculate the joint distribution of C_{t+s} and D_{t+s} . For instance, to price a one-period real bond ($D_{t+1} = 1$ and $D_{t+s} = 0$ for $s > 1$) we need the model to calculate $E_t[\beta(C_{t+1}/C_t)^{-\gamma}]$.

However, if the introduction of the asset does not change the process of consumption, then there is no need for explicitly including the asset (in the budget restriction and invest-

ment opportunity set of investors) when solving the model. This is the case when we are using a model with complete markets (the stochastic discount factor is unique in this case; see the references in footnote 2). In such a model, we could always analyse the pricing of any asset, for instance, some really exotic derivative, without solving the model anew.

We are in a completely different situation when we introduce non-redundant assets, that is, when we add a non-trivial asset to a model with incomplete markets. In that case we have solve the model with all the assets explicitly incorporated since their existence is likely to affect the equilibrium consumption process. Of course, specifying an exogenous process for nonstorable endowments circumvents the problem, which may explain why that approach is so popular.

Most (simulation based) studies of imperfect markets have concentrated on idiosyncratic income risk that cannot be insured and on borrowing constraints/transaction costs. The key findings and references are given in Section 2.6.3. Other studies have developed solution methods and analysed a different set of incomplete markets. For instance, Judd and Guu (2001) show that introducing an option on the stock market can lower the equity premium considerably (see also Den Haan (1996) and Marcet and Singleton (1999) for solution methods and some examples).

4 Summary

The first part of this chapter studies the consumption-based asset pricing model with constant relative risk aversion (CRRA) by comparing the model predictions to the moments (means, variances, and correlations) of US consumption growth and returns on broad asset classes.

It is shown that the smooth consumption growth rate observed in the US makes it hard to explain the equity premium, unless the risk aversion is very high (this is the equity premium puzzle discussed in Mehra and Prescott (1985)). However, a very high risk aversion seems somewhat implausible and it creates other problems. In particular, a large risk aversion coefficient makes it hard to reconcile average consumption growth with the average return on riskfree assets (this is the riskfree rate puzzle discussed in Weil (1989)) and the low volatility of real interest rates (or approximations of it).

One reaction to these problems is to say that the consumption-based model may fail to explain the magnitude of the risk premium, but that it at least can be used to understand

qualitative differences in average return. For instance, the model correctly predicts higher returns on equity than on bonds. One of main issues in the first part of the chapter is how far that argument can be extended. Data on Sharpe ratios for different portfolios sorted on industry, dividend-price ratio, size, book-to-market, and geographical location show that the consumption-based model is unable to explain the cross-sectional variation of Sharpe ratios. The basic problems seems to be that there is a lot of dispersion in Sharpe ratios, but most portfolio returns are only weakly correlated with aggregate consumption growth. It is also shown that the consumption-based model fails badly in explaining time-variation of equity premia. Although the covariance of consumption and equity returns changes over time, it does so in completely the wrong way compared with excess returns.

All these problems have spurred interest in developing the consumption-based model by changing the utility function and/or introducing heterogenous agents. Several of these suggestions (Epstein-Zin preferences, habit formation, and idiosyncratic risk) are analysed at the end of the first part of the chapter. It is found that they are quite successful in solving some of the problems, but that several remain unsolved. In particular, none of the suggestions make much headway with the cross-sectional variation of Sharpe ration.

The second part of the chapter is a theoretical analysis of the equilibrium relation between different asset returns and the consumption process. Simple pricing expressions (based on assumptions about lognormal returns and consumption growth) highlights the basic mechanism of (otherwise) fairly complicated pricing problems. Among the results, we find the following.

First, a positive autocorrelation of consumption growth is likely to give low risk premia on long-lived consumption claims and a negatively sloped real yield curve. Second, the properties of the Fisher equation and the risk premium on nominal bonds depend crucially on monetary policy. For instance, shocks to aggregate demand typically move consumption and inflation in the same direction, so the standard recipe for monetary policy is cushion all such shocks. However, this is likely to increase the inflation risk premium since it makes inflation and consumption more negatively correlated (so the real return on the nominal bond tend to be low when consumption is scarce). Third, options and levered claims will typically have higher risk premia than the underlying asset/unlevered claim because their returns are more volatile.

Simulation models have made fine contributions to the analysis of asset pricing. In particular, they have taught us a great deal about how the properties of the consumption-

based asset model with constant relative risk aversion (CRRA). However, it is clear that the CRRA model has serious problems on many accounts and much of the current research is trying models with different utility functions and/or different measures of consumption. We still know too little about the properties of these new models, for instance, what side effects they have. This is likely to be the next area where simulation models make valuable contributions to the asset pricing literature.

A Data Appendix

The nominal *US stock returns* are from French's (2001) website. These monthly returns are converted to quarterly returns by multiplying the monthly gross returns, for instance, the gross returns for January, February, and March are multiplied to generate a quarterly gross return. The portfolios are formed from NYSE, AMEX, and NASDAQ firms. The aggregate stock market return is a value weighted return of all available returns. The (equally-weighted) portfolios based on industry, size, dividend/price, or book-to-market are from the same data set, but with the firms sorted on the respective characteristic. The ten industry portfolios are for consumer nondurable, consumer durables, oil, chemicals, manufacturing, telecom, utilities, wholesale and retail, financial services, and other. The ten size portfolios are for deciles of firm market values; the D/P portfolios are for deciles of dividend/price; and the B/M portfolios are for deciles of book value/market values.

The dividend-price ratio for S&P500 is from Shiller (2000).

The nominal return on *long US government bonds* is from Ibbotson Associates.

The *international stock and long government bond returns* are from Ibbotson Associates, but come originally from Morgan Stanley and IMF respectively. The international data is for France, Germany, Japan, and United Kingdom, and has been converted into US dollar returns.

Real returns are calculated by dividing the nominal gross return by the gross inflation rate over the same period. Inflation is calculated from the seasonally adjusted CPI for all urban consumers (available at <http://www.stls.frb.org/fred/>).

Quarterly growth of *real consumption per capita* of nondurables and services is calculated from the seasonally adjusted number in NIPA Table 8.7 (available at <http://www.bea.doc.gov/bea/dn1.htm>). The growth rate is calculated as a weighted average of the growth rate of nondurables and the growth rate of services (chained 1996 dollars), where the

(time-varying) weight is the relative (current dollar) size of nondurables in relation to services.

The annualized k -quarter excess return used in the discussion of stock returns predictability are calculated as $R_t^e(k) = (R_t R_{t-1} \dots R_{t-k+1})^{4/k} - (R_{f,t} R_{f,t-1} \dots R_{f,t-k+1})^{4/k}$, where R_t is the quarterly gross return of the aggregate US stock market (see above) and $R_{f,t}$ is the quarterly gross return on a short T-bill.

B Econometrics Appendix

The tests in Table 1 are based on t-tests from a GMM estimation and the delta method.

Testing of the mean, standard deviation, and Sharpe ratio is done in the following way. First, moment conditions for the mean (μ) and second moment (χ) of the random variable x_t are defined as

$$\sum_{t=1}^T m_t(\beta)/T = \mathbf{0}_{2 \times 1} \text{ where } m_t = \begin{bmatrix} x_t - \mu \\ x_t^2 - \chi \end{bmatrix} \text{ and } \beta = \begin{bmatrix} \mu \\ \chi \end{bmatrix}.$$

GMM estimation is trivial and the estimator is asymptotically normally distributed around the true value β_0 as

$$\sqrt{T}(\hat{\beta} - \beta_0) \xrightarrow{d} N(\mathbf{0}_{2 \times 1}, S_0),$$

where S_0 is the variance-covariance matrix of $\sum_{t=1}^T m_t(\beta_0)/\sqrt{T}$, which can be consistently estimated by using the point estimate of β instead of β_0 and then applying the Newey and West (1987) estimator. Second, the standard deviation and Sharpe ratio are calculated as functions of β

$$\begin{bmatrix} \sigma(x) \\ E(x)/\sigma(x) \end{bmatrix} = g(\beta) = \begin{bmatrix} (\chi - \mu^2)^{1/2} \\ \frac{\mu}{(\chi - \mu^2)^{1/2}} \end{bmatrix}, \text{ so } \frac{\partial g(\chi, \mu)}{\partial \beta'} = \begin{bmatrix} \frac{-\mu}{(\chi - \mu^2)^{1/2}} & \frac{1}{2(\chi - \mu^2)^{1/2}} \\ \frac{\chi}{(\chi - \mu^2)^{3/2}} & \frac{-\mu}{2(\chi - \mu^2)^{3/2}} \end{bmatrix}.$$

From the delta method we then have

$$\sqrt{T}[g(\hat{\beta}) - g(\beta_0)] \xrightarrow{d} N\left(0, \frac{\partial g(\beta_0)}{\partial \beta'} S_0 \frac{\partial g(\beta_0)'}{\partial \beta}\right),$$

which is used to construct asymptotic t-tests.

A similar approach is used for *testing the correlation* of x_t and y_t , $\rho(x_t, y_t)$. For expositional simplicity, assume that both variables have zero means. The variances and

the covariance can then be estimated by the moment conditions

$$\sum_{t=1}^T m_t(\beta)/T = \mathbf{0}_{3 \times 1} \text{ where } m_t = \begin{bmatrix} x_t^2 - \sigma_{xx} \\ y_t^2 - \sigma_{yy} \\ x_t y_t - \sigma_{xy} \end{bmatrix} \text{ and } \beta = \begin{bmatrix} \sigma_{xx} \\ \sigma_{yy} \\ \sigma_{xy} \end{bmatrix}.$$

The covariance matrix of these estimators is estimated as before. The correlation is a simple function of these parameters

$$\rho(x, y) = g(\beta) = \frac{\sigma_{xy}}{\sigma_{xx}^{1/2} \sigma_{yy}^{1/2}}, \text{ so } \frac{\partial g(\chi, \mu)}{\partial \beta'} = \begin{bmatrix} -\frac{1}{2} \frac{\sigma_{xy}}{\sigma_{xx}^{3/2} \sigma_{yy}^{1/2}} & -\frac{1}{2} \frac{\sigma_{xy}}{\sigma_{xx}^{1/2} \sigma_{yy}^{3/2}} & \frac{1}{\sigma_{xx}^{1/2} \sigma_{yy}^{1/2}} \end{bmatrix}.$$

The delta method formula (see above) is the applied to this case as well.

References

- Abel, A. B., 1999, “Risk Premia and Term Premia in General Equilibrium,” *Journal of Monetary Economics*, 43, 3–33.
- Altug, S., and P. Labadie, 1994, *Dynamic Choice and Asset Markets*, Academic Press.
- Backus, D., S. Foresi, and S. Zin, 1998, “Arbitrage opportunities in arbitrage-free models of bond pricing,” *Journal of Business and Economic Statistics*, 16, 13–26.
- Backus, D., A. Gregory, and S. Zin, 1989, “Risk Premiums in the Term Structure; Evidence from Artificial Economies,” *Journal of Monetary Economics*, 24, 371–399.
- Bansal, R., and C. Lundblad, 2000, “Fundamental Values and Asset Returns in Global Equity Markets,” manuscript, <http://www.business.wm.edu/conference/Lundblad1.pdf>.
- Bidarkota, P. V., and J. H. McCulloch, 2000, “Consumption Asset Pricing with Stable Shocks - Exploring a Solution and Its Implications for the Equity Premium Puzzle,” mimeo, Kansas State University.
- Black, F., and M. Scholes, 1973, “The Pricing of Options and Corporate Liabilities,” *Journal of Political Economy*, 81, 637–659.
- Bossaert, P., 2002, *The Paradox of Asset Pricing*, Princeton University Press.

- Breeden, D. T., 1979, "An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities," *Journal of Financial Economics*, 7, 265–96.
- Campbell, J. Y., 1986, "Bond and Stock Returns in a Simple Exchange Model," *Quarterly Journal of Economics*, 101, 785–804.
- Campbell, J. Y., 1993, "Intertemporal Asset Pricing without Consumption Data," *American Economic Review*, 83, 487–512.
- Campbell, J. Y., 2001, "Consumption-Based Asset Pricing," in George Constantinides, Milton Harris, and Rene Stultz (ed.), *Handbook of the Economics of Finance*, North-Holland, Amsterdam.
- Campbell, J. Y., and J. H. Cochrane, 1999, "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107, 205–251.
- Campbell, J. Y., A. W. Lo, and A. C. MacKinlay, 1997, *The Econometrics of Financial Markets*, Princeton University Press, Princeton, New Jersey.
- Campbell, J. Y., and L. M. Viceira, 2002, *Strategic Asset Allocation: Portfolio Choice of Long-Term Investors*, Oxford University Press.
- Clarida, R., J. Galí, and M. Gertler, 1999, "The Science of Monetary Policy: A New Keynesian Perspective," *Journal of Economic Literature*, 37, 1661–1707.
- Cochrane, J. H., 2001, *Asset Pricing*, Princeton University Press, Princeton, New Jersey.
- Cochrane, J. H., and M. Piazzesi, 2002a, "Bond Risk Premia," mimeo, University of Chicago.
- Cochrane, J. H., and M. Piazzesi, 2002b, "The Fed and Interest Rates—a High-Frequency Identification," *AER Papers and Proceedings*, 92, 90–95.
- Cogley, T., 1998, "Idiosyncratic Risk and Equity Premium: Evidence from the Consumer Expenditure Survey," mimeo, Federal Reserve Bank of San Francisco.
- Cogley, T., and J. M. Nason, 1993, "Impulse Dynamics and Propagation Mechanisms in Real Business Cycle Model," *Economics Letters*, 43, 77–81.

- Constantinides, G. M., and D. Duffie, 1996, "Asset Pricing with Heterogeneous Consumers," *The Journal of Political Economy*, 104, 219–240.
- Cooley, T. F., and G. D. Hansen, 1989, "The Inflation Tax in a Real Business Cycle Model," *The American Economic Review*, 79, 733–748.
- Cooley, T. F., and G. D. Hansen, 1995, "Money and the Business Cycle," in Thomas F. Cooley (ed.), *Frontiers of Business Cycle Research*, Princeton University Press, Princeton, New Jersey.
- Cox, J. C., J. E. Ingersoll, and S. A. Ross, 1985, "A Theory of the Term Structure of Interest Rates," *Econometrica*, 53, 385–407.
- Dantihine, J.-P., and J. B. Donaldson, 2002, "Labour Relations and Asset Returns," *Review of Economic Studies*, 69, 41–64.
- Den Haan, W. J., 1996, "Heterogeneity, Aggregate Uncertainty, and the Short-Term Interest Rate," *Journal of Business and Economic Statistics*, 14, 399–411.
- Den Haan, W. J., 2001, "The Importance of the Number of Different Agents in a Heterogeneous Asset-Pricing Model," *Journal of Economic Dynamics and Control*, 25, 721–746.
- Epstein, L. G., and S. E. Zin, 1989a, "Substitution, Risk Aversion, and the Temporal Behavior of Asset Returns: A Theoretical Framework," *Econometrica*, 57, 937–969.
- Epstein, L. G., and S. E. Zin, 1989b, "Substitution, Risk Aversion, and the Temporal Behavior of Asset Returns: An Empirical Analysis," *Journal of Political Economy*, 99, 263–286.
- Fama, E. F., and K. R. French, 1988a, "Dividend Yields and Expected Stock Returns," *Journal of Financial Economics*, 22, 3–25.
- Fama, E. F., and K. R. French, 1988b, "Permanent and Temporary Components of Stock Prices," *Journal of Political Economy*, 96, 246–273.
- Fama, E. F., and K. R. French, 1993, "Common Risk Factors in the Returns on Stocks and Bonds," *Journal of Financial Economics*, 33, 3–56.

- Ferson, W. E., 1995, "Theory and Empirical Testing of Asset Pricing Models," in Robert A. Jarrow, Vojislav Maksimovic, and William T. Ziemba (ed.), *Handbooks in Operations Research and Management Science* . pp. 145–200, North-Holland.
- French, K. R., 2001, "U.S. Research Returns Data," <http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/>.
- Giovannini, A., and P. Labadie, 1991, "Asset Prices and Interest Rates in Cash-in-Advance Models," *Journal of Political Economy*, 99, 1215–1251.
- Hansen, L. P., and R. Jagannathan, 1991, "Implications of Security Market Data for Models of Dynamic Economies," *Journal of Political Economy*, 99, 225–262.
- Hansen, L. P., and S. F. Richard, 1987, "The Role of Conditioning Information in Deducing Testable Restrictions Implied by Dynamic Asset Pricing Models," *Econometrica*, 55, 587–613.
- Harrison, M. J., and D. M. Kreps, 1979, "Martingales and Arbitrage in Multiperiod Securities Markets," *Journal of Economic Theory*, 20, 381–408.
- He, H., and D. M. Modest, 1995, "Market Frictions and Consumption-Based Asset Pricing," *Journal of Political Economy*, 103, 94–117.
- Heaton, J. C., and D. J. Lucas, 1996, "Evaluating the Effects of Incomplete Markets on Risk-Sharing and Asset Pricing," *Journal of Political Economy*, 104, 443–487.
- Hendershott, P. H., 1998, "Equilibrium Models in Real Estate Research: A Survey," *Journal of Real Estate Literature*, 6, 13–25.
- Huang, C.-F., and R. H. Litzenberger, 1988, *Foundations for Financial Economics*, Elsevier Science Publishing, New York.
- Jermann, U. J., 1998, "Asset Pricing in Production Economies," *Journal of Monetary Economics*, 41, 257–275.
- Johnson, N. L., S. Kotz, and N. Balakrishnan, 1994, *Continuous Univariate Distributions*, Wiley, New York, 2nd edn.

- Judd, K. L., and S.-M. Guu, 2001, "Asymptotic Methods for Asset Market Equilibrium Analysis," *Economic Theory*, 18, 127–157.
- Kollman, R., 2001, "The Exchange Rate in a Dynamic-Optimizing Business Cycle Model with Nominal Rigidities: A Quantitative Investigation," *Journal of International Economics*, 55, 243–262.
- Labadie, P., 1989, "Stochastic Inflation and the Equity Premium," *Journal of Monetary Economics*, 24, 277–298.
- LeRoy, S. F., and J. Werner, 2001, *Principles of Financial Economics*, Cambridge University Press.
- Lettau, M., 2000, "Inspecting the Mechanism: Closed-form Solutions for Asset Prices in Real Business Cycle Models," mimeo, Federal Reserve Bank of New York.
- Lettau, M., 2002, "Idiosyncratic Risk and Volatility Bounds, or Can Models with Idiosyncratic Risk Solve the Equity Premium Puzzle," *Review of Economics and Statistics*, forthcoming.
- Lettau, M., and H. Uhlig, 2000, "Can Habit Formation Be Reconciled with Business Cycle Facts," *Review of Economic Dynamics*, 3, 79–99.
- Lettau, M., and H. Uhlig, 2002, "Sharpe Ratios and Preferences: An Analytical Approach," *Macroeconomic Dynamics*, forthcoming.
- Lintner, J., 1965, "The Valuation of Risky Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets," *Review of Economics and Statistics*, 47, 13–37.
- Ljungqvist, L., and H. Uhlig, 2000, "Tax Policy and Aggregate Demand Management Under Catching Up with the Joneses," *American Economic Review*, 90, 356–366.
- Lucas, R. E., 1978, "Asset Prices in an Exchange Economy," *Econometrica*, 46, 1426–14.
- Luttmer, E. G. J., 1996, "Asset pricing in Economies with Frictions," *Econometrica*, 64, 1439–1467.

- Mankiw, G. N., 1986, "The Equity Premium and the Concentration of Aggregate Shocks," *Journal of Financial Economics*, 17, 211–219.
- Marcet, A., and K. J. Singleton, 1999, "Equilibrium Asset Prices and Savings of Heterogeneous Agents in the Presence of Incomplete Markets and Portfolio Constraints," *Macroeconomic Dynamics*, 3, 243–277.
- Mehra, R., and E. Prescott, 1985, "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 15, 145–161.
- Merton, R. C., 1973, "Rational Theory of Option Pricing," *Bell Journal of Economics and Management Science*, 4, 141–183.
- Mossin, J., 1968, "Equilibrium in a Capital Asset Market," *Econometrica*, 35, 768–783.
- Newey, W. K., and K. D. West, 1987, "A Simple Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55, 703–708.
- Reffett, K., and F. Schorfheide, 2000, "Evaluating Asset Pricing Implications of DSGE Models," mimeo, Arizona State University.
- Ross, S. A., 1978, "A Simple Approach to the Valuation of Risky Streams," *Journal of Business*, 51, 453–475.
- Rouwenhorst, K., 1995, "Asset Pricing Implications of Equilibrium Business Cycle Models," in Thomas F. Cooley (ed.), *Frontiers of Business Cycle Research*, Princeton University Press, Princeton, New Jersey.
- Rudebusch, G. D., 1998, "Do measures of monetary policy in a VAR make sense?," *International Economic Review*, 39, 907–931.
- Salyer, K. D., 1998, "Crash States and the Equity Premium: Solving One Puzzle Raises Another," *Journal of Economic Dynamics and Control*, 22, 955–965.
- Sharpe, W. F., 1964, "Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk," *Journal of Finance*, 19, 425–442.
- Shiller, R. J., 2000, *Irrational Exuberance*, Princeton University Press, data available at <http://www.econ.yale.edu/shiller/data.htm>.

- Smith, P. N., and M. R. Wickens, 2002, "Asset Pricing with Observable Stochastic Discount Factors," Discussion Paper No. 2002/03, University of York.
- Söderlind, P., 1999, "Solution and Estimation of RE Macromodels with Optimal Policy," *European Economic Review*, 43, 813–823.
- Söderlind, P., 2001a, "Monetary Policy and Bond Option Pricing in an Analytical RBC Model," Working Paper 447, Stockholm School of Economics.
- Söderlind, P., 2001b, "Monetary Policy and the Fisher Effect," *Journal of Policy Modeling*, 23, 491–495.
- Söderlind, P., and L. E. O. Svensson, 1997, "New Techniques to Extract Market Expectations from Financial Instruments," *Journal of Monetary Economics*, 40, 383–420.
- Svensson, L. E. O., 1989, "Portfolio Choice with Non-expected Utility in Continuous Time," *Economics Letters*, 40, 313–317.
- Telmer, C. I., 1993, "Asset Pricing Puzzles and Incomplete Markets," *Journal of Finance*, 48, 1803–1832.
- Vasicek, O. A., 1977, "An Equilibrium Characterization of the Term Structure," *Journal of Financial Economics*, 5, 177–188.
- Vayanos, D., 1998, "Transaction Costs and Asset Prices: A Dynamic Equilibrium Model," *Review of Financial Studies*, 11, 1–58.
- Weil, P., 1989, "The Equity Premium Puzzle and the Risk-Free Rate Puzzle," *Journal of Monetary Economics*, 24, 401–421.
- Zimmermann, C., 1999, "International Business Cycles and Exchange Rates," *Review of International Economics*, 74, 682–698.