CHAPTER 4

THE EMPIRICAL RESULTS

This chapter presents tests of some of the propositions developed in chapter 2. The consumption–based asset pricing model links expected real consumption growth to expected real interest rates. At time t, only the nominal interest rates are available. In order to calculate the expected real interest rates, a model of expected inflation is necessary. The next section investigates the inflation process and calculates expected real interest rates.

4.1 Estimating Expected Inflation

Three different inflation forecasting models are examined in this section. The variable of interest is the rate of change in the price deflator for consumption of non-durables and services. The simplest method presented is the random walk in the inflation rate. Next a univariate time series representation is estimated for the entire series and a moving sample window. Finally, the methodology of Fama and Gibbons (1984) is used to calculate a Treasury bill based model of inflation. Forecasts of inflation from t to $t+1$ for all of the models are based on information available at time t. The parameters of each model are re-estimated at every point in the series and the j–step ahead forecasts are calculated at every point.

4.1.1 Quarterly Data

The quarterly inflation process over the 1948:2–1985:2 period[1] appears to follow an IMA(1,1). An examination of the autocorrelogram shows a very slow decay, a cut through the zero axis and a movement to significant negative values at high orders. This type of structure is typical of mean non-stationary processes. The other candidate process is an ARMA(1,1). This process is less likely to be the correct process because of the shape of the autocorrelogram. When the models are fit,[2] the AR parameter of the mean reverting model is less than two standard errors from one. The residual sum of squares of the IMA(1,1) and the ARMA(1,1) are virtually the same – largely because of the closeness of the AR parameter to one. This makes it difficult to choose the best model. A Dickey–Fuller unit root test[3] does not provide evidence against the null hypothesis of a unit root. The IMA(1,1) is chosen on the basis of the initial identification. The estimates of the IMA(1,1) model appear in table 1.

Inflation forecasts are obtained in two ways. First, the IMA(1,1) is estimated over an initial period, 1948:2 to 1958:1. The fitted values provide the inflation forecasts in this period. After this initial estimation period, the model is re-estimated at each point in the time series and j–step ahead forecasts are obtained. The second estimation strategy involves a moving sample window. After 1962:1,

[1] There is data available for the 1946:2–1948:1 period. This part of the sample is excluded because of the interventions in the time series that result from the lifting of war–time price controls.

[2] All of the time series models are estimated with the exact likelihood algorithm available in the SCA computer package.

[3] The Dickey–Fuller likelihood ratio test statistic for unit roots in series Y_t is calculated from $\mathrm{H}_0 : (\alpha, \beta, \rho) = (\alpha, 0, 1)$ in $Y_t = \alpha + \beta(time) + \rho Y_{t-1} + \epsilon_t$. They provide an empirical distribution for the statistic in their 1981 paper. However, they admit that their test may lack power.

when a point is added to the sample, a point at the beginning of the sample is dropped. The length of the window is 60 quarters.[4] These estimates will be referred to as *Time Series Window*. As with the first method, the parameters of the model are re-estimated at every point in the time series. Figure 1 provides graphs of the moving average parameter estimates over time. Plots of the actual and fitted values from the two models are shown in the appendix.

The final model considered follows Fama and Gibbons (1984). In this representation, a time series model is applied to the *ex post* real interest rate. Forecasts are calculated and subtracted from the yield on a current Treasury bill to get an implied measure of expected inflation. The identification of the time series model for the real rate is similar to the exercise for the inflation rate. Two *ex post* real rate series are considered: average rate on a 90 day Treasury bill over the quarter and a spot rate in the middle month of the quarter. Both series are very similar. The average series has a larger moving average component by construction. The *ex post* real rate series appears to be mean non-stationary but the evidence is not as clear as with the inflation rate.[5] Again two models were fit: an IMA(1,1) and an ARMA(1,1). The IMA(1,1) model was chosen. The estimates of the IMA(1,1) parameters for the entire sample appear in table 1.

As with the time series model on the inflation rate, the univariate representation of the *ex post* real rate was estimated at each point in time after an initial estimation period. Figure 2 shows how the moving average parameters of the model applied to the spot and average data vary through time. As expected, the

---

[4] The window is chosen to be long enough to allow the parameter values to exhibit some stability. If the window was shorter, say 20 quarters, the model identification might change as a result of a different location in the business cycle.

[5] The Dickey-Fuller unit root test does not provide much evidence against the null hypothesis of a unit root.

moving average parameter associated with the average interest rate data has a higher value. But it is remarkable how similar the parameter sequences are. It is also interesting to note the jump downward in the parameters in the 1973–1974 period and the 1979–1981 period. These periods coincide with recessions. The decline in the value of the moving average parameters indicate that a greater proportion of change in the real rate is unexpected. The inflation forecasts generated from these models will be referred to as *Treasury Bill (Avg.)* and *Treasury Bill (Spot)*.

4.1.2 Annual Data

While annual inflation data exist from 1872–1984, the one-year yield series is available only from 1900 to 1984. The time series based model of the inflation rate uses the entire series but the forecasts are restricted to the 1900–1984 period.

The identification of the annual inflation process is more straight-forward than the quarterly inflation process. The annual inflation rate appears to follow an ARMA(1,1). Estimates of the parameters of this model for the full sample are provided in table 1. One step ahead forecasts are generated by re-estimating the parameters of the model at each point in the series starting in 1900. As with the quarterly data, a time series window model is also estimated. The window length is 20 years. This is a longer period than the quarterly model. It is difficult to have much confidence in the time series estimates if fewer than 20 data points are used. Forecasts are generated from 1900 to 1984. Figures 3 and 4 show the time series behavior of the ARMA parameter estimates. The both the MA and AR parameters on the time series model are fairly stable through the sample. The time series window estimates track these parameters closely until 1950 after which there is often a wide divergence between the windowed and unwindowed estimates.[6]

The final model considered is a time series model for the *ex post* real rate. In the 1900–1983 period, the realized real rate appears to follow an ARMA(1,1). Since the sample size is smaller, the fitted values are used for forecasts in the 1900–1919 period. The model is re-estimated at every point in the time series after 1919 and one-step ahead forecasts are obtained. The full sample estimates of this model are provided in table 1. Figure 5 plots the estimates of the ARMA

[6] This could be due to the small size of the window.

parameters through time. The parameter estimates are reasonably stable except during the 1929–1934 period. In the depression, the AR parameter moves close to one and the MA parameter drops toward zero. This suggests that the process resembles a random walk in this particular sample. As in the quarterly sample, during recessionary periods changes in the real rate appear to be unpredictable.

4.1.3 Forecast Evaluation

Table 2 provides some summary statistics for evaluating the the inflation forecasting models. Statistics are provided for the forecast horizons of one to four periods for the quarterly data and one year for the annual data. The span of the data corresponds to the availability of the financial instruments. Three evaluation criteria appear: root mean squared error (RMSE), mean absolute error (MAE) and the correlation between actual and forecasted.

In all the forecast horizons, the time series model provides the lowest RMSE. In addition, the time series model has the lowest MAE in the two, three and four quarter ahead horizons and in the one year (annual) ahead horizon. This evidence suggests that the time series based model should be used as the model of inflation.

The performance of the time series model of inflation seems to run counter to the results of Fama and Gibbons (1984) who document the superiority of the Treasury bill model. But on close examination, the results are fully consistent even though quarterly rather than monthly data are being used and the personal consumption deflator rather than the CPI is being employed. In figure 6, the cumulative RMSE and MAE for the one quarter ahead forecasts are presented.[7] Through most of the sample, the Treasury bill model provides the lowest RMSE but there is a substantial deterioration of the forecast power in 1979–1982 period. Fama and Gibbons document a similar deterioration in their final sub-period, 1977-1981. Similar graphs for the other horizons are found in appendix E. For the two to four quarter horizons, a similar deterioration in the forecast power in this period is found – but the dominance of the Treasury bill model in the earlier period is not as evident.

[7] Figure 15 in appendix E provides graphs of the actual forecast errors.

In figure 7, the cumulative RMSE and MAE for the annual forecasts are presented. There is a large jump upwards in the forecasting errors in the time of depression. The quarterly data exhibited jumps in the errors in the 1973–1975 and 1979–1982 periods. The annual data show a deterioration in the forecast power in the 1929–1934 period but there is no deterioration in the 1970's. The correlations between actual and fitted are lower for the annual data indicating that it is more difficult to forecast one year ahead. Finally, there is not much difference in the forecasting power of the bond based model or the time series based model. In the first 20 years of evaluation, the bond model does better than the time series model. But this could be due to the fact that the bond model is using fitted values during these years whereas the time series is using out-of-sample forecasts.

It is important to note that these models are just approximations of the true expectations generating mechanism. There is no reason to presume that the expectation generation process remains fixed through time. New variables may enter the process and old variables may be dropped. Letting the parameters vary through time may help the time series models adjust to this type of change in the true underlying structure.

While the model evaluation seems to point to the time series representation of inflation, many of the results are replicated using the alternative models of expected inflation. This provides a test of how sensitive the results are to the use of different expected inflation measures.

4.2 Real Interest Rates and Yield Spreads – Preliminary Data Analysis

The inflation forecasts in the previous section are now applied to the nominal interest rate structure in order to calculate the expected real rates. The theory suggests that expected real rates should be related to consumption growth. The preliminary data analysis provides an informal way to assess the similarities and differences in the time series. Tables 3–5 present some summary statistics for quarterly and annual measures of real consumption growth, expected real interest rates and yield spreads. These statistics are based on average nominal interest rates. Tables 25 and 26 in appendix D replicate these statistics using the spot interest rates for the quarterly data. These tables also provide preliminary data analysis for the real rates calculated with different models of inflation: time series, time series window, Treasury bill and random walk. The *ex post* real rate is also included.

4.2.1 Quarterly Data

Summary statistics are provided in table 3 for the expected real rates and growth in consumption of non-durables and services. The results of Working (1960) show that the difference between two series that are averaged will induce a first-order moving average process in the differenced series.[8] If consumption follows a geometric random walk, then the growth in consumption should follow a moving-average process with a first-order autocorrelation coefficient of .25 and higher order autocorrelation coefficients of zero. The results in table 3 show that the sample first-order autocorrelation coefficient for the one quarter consumption

[8] Both Hall (1985) and Breeden, Gibbons and Litzenberger (1985) have applied these results to the consumption data.

growth is .245 which is close to the hypothesized value but the higher order coefficients are not all close to zero. Since the consumption growth is a continuous–time growth rate, the longer growth rates are just sums of the one-quarter rates. Hence, the two to four quarter growth rates should follow two to four quarter moving-average processes. Furthermore, the first-order serial correlation coefficient will increase as more series are summed.[9] This is more or less what is seen in table 3.[10]

An examination of the autocorrelogram of the real rate process reveals a different picture. As mentioned earlier, the real rate process has a slow decay that is either indicative of a non-stationary process or a mean reverting process with an AR coefficient close to unity. The two to four quarter expected real rates are not sums of the one quarter rate as was the case with consumption growth. An across horizon comparison suggests that all of the expected interest rates have a similar autocorrelation function.

Table 4 presents the preliminary data analysis for the expected real yield spreads and real consumption growth. The consumption growth statistics are again displayed on the top line of each panel. The expected j–period real yield spread measure has a different time series behavior than the expected real interest rates which are its components. The autocorrelograms of the three *ex ante* yield spreads show a sharp decay to zero. The autocorrelation structure is roughly similar to that of the consumption growth measures.

---

[9] If $\Delta C_t = (1 - \theta L)\epsilon_t$, where L is the lag operator, then the sum of n adjacent series is $\sum_{j=1}^{n} \Delta C_{t+1-j} = (1 - \sum_{j=1}^{n-1}(\theta - 1)L^j - \theta L^n)\epsilon_t$. From this expression, the autocorrelations of the summed series can be derived. For example, if in the MA(1) ρ_1= .25, then MA(2):ρ_1=.60, MA(3):ρ_1=.77 and MA(4):ρ_1=.82.

[10] Note that the length of the series is different across horizons. Also, the four quarter growth rates are not the sum of the one quarter growth rates. They are the sum of the one quarter not seasonally adjusted growth rates.

The summary statistics suggest that the expected real rates are either mean non-stationary or mean reverting with a high autoregressive coefficient yet the spread and the consumption growth series are clearly mean stationary. This is not necessarily inconsistent with the model. The growth in consumption should equal the expected real rate plus noise. If noise is added to the expected interest rate process, it should end up looking random like the consumption growth process.

A more detailed look at the quarterly data is presented in appendix tables 27–30. These tables present measures of skewness and kurtosis as well as two normality test statistics. Given that the data are stationary, these test statistics provide information about the distribution of the data. Two normality test statistics are the studentized range and the Kolmogorov D–statistic.[11] The D–statistics show that we cannot reject at the one percent level of significance the null hypothesis of normality for all of the consumption growth series and for the two and three quarter yield spread series. There is evidence against the null hypothesis with the other interest rate measures. Joint lognormality of the marginal rates of substitution and the interest rate measures is important to the linear specification of the model. The normality test results suggest that this assumption is not supported by the data when applied to the real interest rates. But the assumption is supported for the yield spread specification. Additional tests presented in

[11] Order the sequence $\{x_j : 1 \leq j \leq n\}$ such that $\{x_1 \leq x_2 \leq \ldots \leq x_n\}$. The Kolmogorov D–statistic is defined as $D = \sqrt{n}\,\sup|F_n(x) - F(x)|$ where $F_n(x)$ is the observed cumulative distribution function and $F(x)$ is the uniform cumulative distribution. The D–statistic measures the greatest divergence between the two distributions. Given a null hypothesis that the sequence is normally distributed, a high value (close to unity) provides evidence against the null hypothesis. However, the test assumes that $\{x_j\}$ are independently identically distributed random variables. Some caution should be exercised when applying this test to a time series, such as consumption growth, where it is difficult to argue that the drawings are i.i.d.

section 4.3 suggest that these results could be due to a non-constant variance in the expected real rate process. The variance in the expected yield spread process is roughly constant.

Some of the data are displayed in figures 8 and 9. If the assumptions of the linear specification are accurate, then the consumption growth series should look like the expected interest rate series plus noise. If the noise component is small then, the series should exhibit similar co-movement. Figure 8 shows the one to four quarter consumption growth plotted against the one to four quarter real rate (calculated with the time series model of inflation). The interest rate variables are lagged to match the span of the consumption growth, i.e., a bill that has three quarters to maturity at time t is matched with consumption growth from t to $t+3$. If these series exactly coincide, then the expected real interest rate perfectly forecasts future consumption growth. The plots show that the consumption series have constant variances. However, the expected real interest rates appears to have non-constant variances. Second, the consumption and interest rate series move with similar long-term trends; both of the series would be quite similar if a high-order moving average were applied to the data. Third, the interest rate deviates sharply from the consumption growth series in the 1979–1982 period. This coincides with the period when the Federal Reserve announced it was changing the focus of its monetary policy. This period is characterized by large unprecedented swings in interest rates.

The plots of the yield spread and consumption growth are presented in figure 9. There are a number of interesting features to these graphs. First, the variance of the spread series seem more constant than the variance of the expected real interest rate process. This is especially evident with the two and three quarter spreads. Second, the deviation in 1979–1982 largely disappears. Third, in the 1964–1967 period, the spread measure is flat whereas there is some volatility in the consumption measure. It is possible that consumers made forecasting errors during this period or the inflation forecasting model used in the analysis does not adequately capture consumers' expectations. The fourth and perhaps most striking point that is evident from figure 9 is the similarity in the movement of the two series. Examination of the figure suggests that the yield spread may contain information relevant for the forecasting of consumption growth.

4.2.2 Annual Data

Table 5 presents summary statistics for the annual data. The top panel of this table considers the 1900–1984 period. The real interest rate variable is a one year corporate bond yield. The spread variable is constructed by taking the natural logarithm of the ratio of the one year corporate yield and the yield on 30–60 day Commercial paper (1900–1919) and 90 day Treasury Bills (1920–1984). The second panel shows the annual data for the 1953–1984 period. The interest rate variable in this sub-period is the rate on a one year government bond. The spread is constructed with the yields on two government instruments. This sample is the more reliable than the longer period because the government instruments are available and because this sub-period follows the Treasury–Federal Reserve Accord.

In the 1900–1984 period, the consumption growth variable appears to be mean stationary with the autocorrelation function trailing off to zero by the third lag. The first-order autocorrelation coefficient exceeds the .25 value that is expected if the underlying series follows a geometric random walk. Caution should be taken in interpreting this time series due to the splicing of the data in 1929. The 1900–1929 data are based on interpolated values from five year sampling intervals. This induces a higher order moving average process in the early data and high autocorrelations are expected. In the 1953–1984 sample, the sample first-order autocorrelation coefficient is .251 which is more in line with the quarterly results.

While the autocorrelation function of the expected real rate in the 1900-1984 period presented in table 5 looks similar to the expected real rate in the quarterly sample documented in table 3, the annual spread measure takes much longer to decay to zero. This could be due to the splicing of instruments in the short term yield series or due to the construction the spread as the difference between a corporate bond and a government bill. The 1953–1984 sub-period is chosen to avoid both of these problems and the autocorrelations are very similar to the quarterly results.

A more detailed examination of these series is contained in appendix table 31. The evidence suggests that the null hypothesis of normality cannot be rejected at the one percent level of significance for the yield spread measure in the full sample and the consumption growth measure in the 1953–1984 period. Normality can be rejected for the real interest rate in both samples. These results may cause problems with the linear specification – at least when it is tested using real interest rates.

Graphs of the data appear in figures 10 and 11. As with the quarterly plots, the data are arranged so that if the series exactly coincide, the interest rate variable provides a perfect forecast of the growth in the economy. The one year growth in consumption and the one year expected real interest rate for the 1900–1984 period are plotted in figure 10. There is a large divergence in the two series during the great depression. Consumption growth is low whereas the expected real yield on the bond is high. The nominal yield on the bond is relatively low. During this time there is a considerable deflation. The time series based forecasting model predicts continued deflation. This leads to the high expected real yields. It is possible that the true inflationary expectations are different during this period. Examination of the inflation forecasting errors of the time series model show large errors in this period. It is also possible that the real expected yields are high because the default risk for the corporate bonds increases. The theory is presented in terms of default-free instruments. The expected real rate roughly tracks the consumption growth

– especially after 1935. There is another large deviation after 1979. Figure 10 also highlights the 1953–1984 period. The change after 1979 is particularly evident in this figure. The deviation coincides with the Federal Reserve's announcement that it would change the focus of its monetary policy.

Plots of the yield spread measures are in figure 11. The divergence of the two series during the depression is not as wide as seen with the expected real interest rate. The expected spread series roughly moves with consumption growth through much of the sample. The relation seems particularly strong after 1950. The graphs suggest that there may be important information in the expected spread that could be used to help forecast growth in consumption.

4.3 The Regression Results: Linear Specification

If the consumption-based model is an adequate description of the world and if the joint distribution of the consumption growth and real rate series is stationary lognormally distributed, then it may be possible to extract information about future consumption growth from the real rate with linear least squares. The preliminary data analysis indicates that we are less likely to find information in individual interest rates. The summary statistics and the graphs suggest that the spread specification better fits the assumptions of the linearized model.

4.3.1 Quarterly Data

The regression results in table 6 for expected real interest rates show that the interest variables appear to have no power to explain variation in future consumption growth. None of the coefficients on the interest rate variables are more than one standard error from zero. All of the standard errors are corrected for a moving average induced in the residuals as a result of overlapping dependent variables. Standard errors are also adjusted for conditional heteroskedasticity.[12]

There may be econometric problems associated with these regressions. The first is the errors in the variables problem. The expected real rate is forecasted and subject to error. As a result, the parameter estimates could be inconsistent. The second problem relates to the assumption of stationary joint lognormality. If

[12] This procedure follows White (1980) and Hansen (1982). Let j be the order of the moving average process. Define $S_W = \frac{1}{T}\sum_{i=-j}^{j}\sum_{t=1}^{T} u_t' u_{t-i} X_t' X_{t-i}$, where u_t are the regression errors. The variance covariance matrix is then $\frac{1}{T}(X'X)^{-1}S_W^{-1}(X'X)^{-1}\frac{1}{T}$. This matrix follows the variance covariance matrix presented Theorem 1 of White (1980) if $j = 0$ and Theorem 3.1 of Hansen (1982) for the more general case of $j \geq 0$.

this assumption is violated, then the regressions are not well specified.

The errors in the variables problem can be addressed with the use of instrumental variables estimation. Table 7 presents some results[13] using a linear version of the GMM technique of Hansen (1982). The instrumentation consists of a constant, the expected real rate from a time series model[14] on the realized real rate and the log of the ratio of yields on Moody's BAA and AAA rated bonds.[15] The instrumental variables are all important predictors of the realized real rate. The GMM estimates of the risk aversion parameter are all positive but not different from zero at conventional levels of significance. This contrasts with the negative point estimates of the risk aversion parameter in two of the four interest rates in the OLS results. The GMM estimation does not provide evidence against the specification with one over-identifying restriction – but there does not seem to be much explanatory power in the specification.

Figure 12 provides a plot of a three year moving window estimate of the variances of joint processes. There is a large increase in the variance in the 1979–1982 period. The linear specification assumes that the variance of the joint process is constant and it is grouped into the intercept term for estimation. The plot of the variance suggests that this assumption is violated. The graph suggests that a closer approximation might be obtained by using the difference in the variance of two processes. This is precisely what the yield spread specification groups into the intercept term.

---

[13] Estimates were obtained using a modified version of a GMM program written in GAUSS by David Runkle and Gregory Leonard.

[14] Since the time series model is re-estimated at each point in time, this variable is predetermined and hence a legitimate instrument.

[15] Keim and Stambaugh (1986) present evidence that this variable is able to predict excess bond returns.

The OLS estimates of the yield spread specification are presented in table 8. Following equation (2.17), both the expected yield spread and the level of the short-term rate are included as explanatory variables. The full sample results suggest that this specification has only marginal explanatory power: six percent of the variation in the two quarter consumption growth is explained but none of the variation in the one and three quarter consumption growth is explained. Some sub-period results are also presented. The first sub–period (ending in 1971:4) is similar to the overall period where there is only slight explanatory power. The second sub-period shows a different picture. The three yield spreads can explain 3%, 23% and 24% of the variation in one to three quarter consumption growth.

The instrumental variables estimates for the full sample are presented in table 9. As with the GMM results in table 7, none of the parameters are significantly different from zero. However, the point estimates are close to those delivered by OLS. The OLS parameter estimates on the one to three quarter yield spread are .28, .86 and .42 while the GMM estimates are .11, .76 and .21. As mentioned earlier, if the estimates are fairly close, then it is less problematic to use the OLS for prediction.

The regression results show that there is no power to predict consumption growth when examining quarterly expected real interest rates. However, there is some ability to predict consumption growth when the yield spread is combined with the short term rate. The explanatory power seems to be concentrated in the second sub-period. A possible explanation for the yield spread having more predictive ability is a changing variance of the joint process. The yield spread specification uses a difference in the variance while the interest rate specification uses the level of the variance.

4.3.2 Annual Data

The results of the regression analysis of the real interest rate and annual consumption growth are in table 10. Similar to the quarterly results, very little information seems to be contained in the level of the real rate. The coefficient on the real rate in the overall period is negative which is the opposite sign that standard economic theory would predict. In the post-depression sample, the coefficient is positive but the standard error is large. As with all the OLS regressions, the standard errors are corrected for a first-order moving average process and for conditional heteroskedasticity.

The GMM estimates presented in table 11 are similar to the OLS estimates. This could indicate that the errors in the variables problem is not that severe in the annual data. But there may be other problems with the linear specification.

Figure 12 presents a plot of five year moving window estimate of the variance of the joint process for the annual data. There is a large jump in the variance in the 1929–1934 period. This is possibly the reason that the interest rate series and the consumption growth series widely diverge in this period.

Table 12 presents the results using the yield spread specification. In the overall period, 1901–1984, neither of the coefficients are significantly different from zero. In the post-depression sample, the yield spread enters with a coefficient that is more than three standard errors from zero. Similarly, the yield spread enters with a coefficient that is greater than two standard errors from zero in the post–1953 sample. These results are corroborated by the GMM results presented in Table 13. The point estimates on all the parameters are slightly larger. The coefficients on the yield spread are positive and in the post-depression sample are 1.6 standard errors from zero.

The annual results suggest that the yield spread variable has some ability to predict consumption growth – but this ability is limited. The real rate has virtually no power to predict. The lack of power could be due to a number of factors such as: errors in the variables, a non-constant variance of the joint process, inflation forecasting errors, government interventions, the averaging of the interest rates, the unavailability of a long-term government instrument before 1953 or a mis-specification of the underlying model. The evidence presented in this section points to large inflation forecasting errors in the depression as well as a non-constant variance as the likely culprits. The reason the spread specification tends to do better is probably because it allows the variance to change through time.

4.3.3 Nominal Interest Rates and Consumption Growth

Most of the econometric problems are caused by the need to estimate a real rate of interest at time t for t to $t+j$. If the assumption of logarithmic utility is maintained, then the first-order conditions can written in terms of nominal rates of interest and nominal consumption growth. Since the j–period nominal interest rate is in the information set at time t, the expectation operator can be dropped from the regressor. Previous studies, such as Hansen and Singleton (1982, 1984), Ferson (1983), Brown and Gibbons (1985), Dunn and Singleton (1985), Hall (1985), Mankiw (1985) and Miron (1985) could not reject a relative risk aversion coefficient of unity.

The nominal specification reverses the steps of the real model. In the nominal specification, the nominal rate is used to forecast nominal consumption growth. The inflation forecasts are then subtracted from the fitted values in the nominal regression to provide forecasts of the growth rate in real consumption. The advantage of this specification is the ability to avoid the errors in the variables problem. Another advantage is that the assumption of *joint* lognormality of the interest rate and consumption growth processes need not be made.

The results of the nominal specification for the quarterly data are provided in tables 6 and 8. Two R^2s are presented. The first is the coefficient of determination resulting from the regression of nominal interest rates on nominal consumption growth. The second R^2 results from the regression of the fitted values of nominal growth model less the inflation forecasts on real consumption growth. In the interest rate specification, the nominal-based predictions do better than the expected real predictions. Table 6 shows that the nominal-based model can explain 2–16% of the variation in real consumption growth while the expected real rate can explain none. The nominal-based predictions also tend to do better in the overall period in the yield spread specification. Table 8 shows that these predictions can account for 5–14% of the variation in the real consumption growth. The nominal based predictions also tend to do better in the first sub-period. The second sub-period is characterized by similar predictive power for both the nominal-based and the expected real-based.

The annual results are presented in tables 9 and 11. As with the quarterly data, table 9 documents the nominal-based predictions consistently outperforming the expected real interest rates in explaining real consumption growth. In the yield spread formulation, the nominal-based and expected real models explain a similar amount of the variation in the real consumption growth.

In addition to avoiding some of the errors in the variables problems, the nominal-based predictions outperforming the expected real interest rate predictions is consistent with the conjecture that the changing variance of the joint process may be causing estimation problems. The nominal-based model does not rely on the assumption of stationary joint lognormality of the real interest rate and consumption growth. The dominance of the nominal-based predictions is not as evident when looking at the yield spread specification. Sometimes the expected yield spread has more power than the nominal-based predictions. This is consistent with the conjecture that the yield spread specification minimizes the changing variance problem by utilizing a difference in two variances.

4.4 GMM Estimation: Non-Linear Specification

This section presents results of the direct estimation of the first-order conditions using the Generalized Method of Moments. This technique has been used to estimate the linear version of the model. The advantage of the non-linear estimation is that the strong distribution assumptions necessary to linearize the model need not be made. Unfortunately, the non-linear estimation does not provide forecasts. But we can learn something about the parameter values from this estimation technique. If the assumptions of the linear version of the model are true, then the non-linear estimation should deliver the same parameter estimates.

4.4.1 Quarterly GMM Estimation

The results of the generalized method of moments estimation of the non-linear first-order conditions are in table 14. The instrumentation consists of a constant, the expected real rate from a univariate time series model, lagged consumption growth and the natural logarithm of the ratio of the yields on Moody's BAA and AAA rated bonds. The point values of the coefficient of relative risk aversion range from .6 to 1.3 in the one to four quarter samples. However, the standard errors are generally large. This range is fairly narrow compared to the values reported in other studies.[16] For example, using monthly data, Hansen and Singleton report estimates from -1.3 to 1.6. The quarterly estimates in table 14 are always positive

[16] Using monthly data, Brown and Gibbons (1985) estimate a range for the coefficient of relative risk aversion of .09 to 7.00. Dunn and Singleton (1985) report values between 1.22 and 1.91 for the single equation estimation and 2.50 to 3.45 for the multiple equation estimates. With quarterly consumption data, Mankiw, Rotemberg and Summers (1985) document values from .09 to .51 in the case where utility is assumed to be separable between consumption and leisure. Mankiw (1985) reports a range between 2.43 and 5.26 using only the fourth quarter data for each year. Using seasonally unadjusted data, Miron (1985) documents a range of .02 to 1.71 for the risk aversion parameter. All of these studies use the GMM technique to obtain parameter estimates.

– the sign predicted by economic theory.

The point estimates of the linear version of the model for real interest rates range from 8.9 to 22.7 as reported in table 7. Again the standard errors are large. The wide divergence from the non-linear estimates seems to suggest that the linear version with real interest rates is mis-specified. The GMM results reported in table 9 using the yield spread specification deliver estimates of 9.1, 1.3 and 4.5. While these estimates are still high compared to the non-linear estimates, they suggest that the yield spread specification may be superior to the interest rate specification. If the true parameter of constant relative risk aversion does fall in the range of .6 to 1.3, then only the yield spread regressions provide that range for the coefficient – particularly in the 1972:1–1985:3 sub-period.

Table 14 also provides a test of the over-identifying restrictions. None of the specifications is rejected at the .01 level of significance. The tests indicate that there does not seem to be much evidence against the constant relative risk aversion (CRRA) specification. Previous research has concentrated on single period returns and the model has received mixed support. Hansen and Singleton (1984) cannot reject the model when a single asset is used. However, they find evidence against the model when more than one asset is used. Dunn and Singleton (1985) using a formulation that allows for durable goods find similar results in testing single equation versus multiple equation specifications. Mankiw, Rotemberg and Summers (1985) find evidence against the model when they use the seasonally adjusted quarterly data in both cases of separable and non-separable utility in consumption and leisure. Yet when they re-estimate using only the fourth quarter data, they do not reject the model. Mankiw's (1985) specification allows for durables as well as non-durables to be incorporated into the model. Using fourth quarter data, he does not find evidence against the specification. Miron (1985) uses seasonally unadjusted data and finds no evidence against the over-identifying restrictions.

The results in Table 14 are not directly comparable to any of the above studies because different asset return data are used and the model is often formulated in different terms in other research, i.e., allowing for consumption services from durables. The results in table 14 indicate that the strongest case against the model can be made from the one quarter rate of change specification. There is little evidence against the model when it is tested with longer term growth rates. A possible explanation for this phenomena is measurement error in the data. Suppose that the measured consumption data sometimes incorporate some information from the next period. This could be caused by the retail sales data collection being off by a few days or the trending of the services data (some of which is collected semi-annually and annually). It could also be result from the forward smoothing of the seasonal adjustment procedure. If this is the case, the measured marginal rate of substitution and the asset return are improperly matched in the first-order conditions and the lagged consumption growth is no longer a valid instrument in the estimation. This is less of a problem when longer term growth rates are used relative size of the error is smaller.

The measurement error explanation is consistent with the pattern of the χ^2 statistics in table 14. The size of the test statistics decreases from the one to three quarter models indicating that there is less evidence against the model as longer growth rates are used. The four quarter growth model reported in table 14 is not directly comparable to the one to three quarter models because the not seasonally adjusted data is used in this case. The size of the measurement error is probably smallest when the model is tested using the seasonally unadjusted data because the series is not forward filtered. The results suggest that there is little evidence against the four quarter specification.

For comparison, the same first-order condition is tested with negative exponential utility[17] which implies constant absolute risk aversion (CARA). In contrast to the CRRA results, the CARA model is rejected at the .05 level in three of the four maturities. These results suggests that the CRRA is a superior specification.

[17] Ferson (1983) provides some tests of the model with a CARA specification using a maximum likelihood estimation technique. There are no studies that have tested the CARA model using the instrumental variables methodology.

4.4.2 Annual GMM Estimation

The generalized method of moments estimation of the non-linear first order conditions is presented in table 15. The instrumentation consists of a constant, the expected real rate generated from a univariate time series model and a lagged consumption value. In the overall sample, the coefficient of relative risk aversion is -5.3. This implies non-concave preferences. The point estimate from the the GMM estimation for the linear version of the model presented in table 11 is similar at -4.7. The standard errors are large in both estimates. These results could be caused by the trended consumption data in the pre-depression sample. The post-depression sample delivers a different set of results. The non-linear estimate of the CRRA parameter is 1.6. This is similar to the parameter found when estimating the model with quarterly data. The GMM estimation for the linear interest rate version of the model reported in table 11 delivers a parameter estimate of 7.4 while the linear yield spread specification reported in table 13 delivers a parameter estimate of 1.01. The linear yield spread point estimate is much closer to the non-linear estimate and again implies that the yield spread model is better specified.

Table 15 also provides a test of the over-identifying restriction. There is no evidence against the model in any of the annual samples. In contrast to the quarterly results, the alternate CARA specification is not rejected at conventional levels of significance although the χ^2 values are always higher than those of the CRRA specification.

The non-linear estimation provides one more piece of evidence in favor of the yield spread rather than the interest rate specification. When the model is estimated the non-linear form, the strong distributional assumptions about the behavior of the joint process need not be made. But if these distributional assumptions are accurate, then the non-linear and linear parameter estimates should be identical – at least in large samples. The evidence suggests that this is not the case with the real interest rate specification. The yield spread specification tends to have parameter estimates that are closer to the non-linear results. The most likely reason for the difference is that the interest rate specification forces the level of the variance of the joint process to be constant. On the other hand, the linear yield spread specification only requires that the difference in the variances of the joint processes remain constant through time.

4.5 Alternative Predictors of Consumption Growth

In the economics literature, there is considerable interest in the predictability of consumption. The Life Cycle–Permanent Income Hypothesis suggests that agents should not alter their consumption plans as a result of shocks in their current income - plans will only be changed if there is a revision in expectations about *permanent* income. Hall (1978) tested this proposition in the context of the first-order condition (2.2). If quadratic utility is imposed and the real interest rate is assumed to be constant, then consumption should follow a random walk. This implies that changes in consumption should be unpredictable. In his empirical tests, Hall found that the consumption process seemed to follow a random walk. However, Hall found that stock returns had some ability to predict consumption. Similarly, Flavin (1981) found that consumption is sensitive to current income even when the role of current income signaling changes in permanent income is taken into account. This provides evidence against the Permanent Income Hypothesis.

Much of the recent research has concentrated on explaining this excess sensitivity. New specifications of the maximand allow for a technology that produces consumption services from durable goods. Liquidity constraints have also been added to the formulation. Recently, Hall (1985) has investigated the role of changes in the real interest rates. In his original formulation and in Flavin's work, the real interest rate is assumed to be constant. If the rate is allowed to vary, then changes in the rate should be related to changes in consumption. As shown above, the coefficient that measures the sensitivity of consumption growth to changes in the expected real rate is the intertemporal elasticity of substitution – or the inverse of the coefficient of relative risk aversion.

Hall's (1985) findings suggest that there is little or no power in the expected real rate to predict consumption growth. This study suggests that there may be problems in testing the relation in terms of one quarter consumption growth and a one quarter expected real interest rate. Indeed, if an alternate interest rate measure such as the yield spread is used, the evidence suggests that there is some power to predict consumption growth.

An interesting question to ask is how important are interest rate variables in predicting consumption growth. Two alternative formulations are suggested. The first includes lagged consumption growth. Hall (1978) found that only the first lag in the consumption level was an important predictor variable. In terms of consumption growth, this implies that no lagged growth terms should be important. The second formulation allows stock returns to predict consumption growth. This variable is a proxy for income changes. The theory suggests that this variable should not have more power to predict consumption growth than the interest rate variable.

Tests are constructed to determine whether there is more information in lagged consumption and real stock returns than in interest rates and yield spreads. Both within-sample and out-of-sample forecasting performance are evaluated. The predictions of the yield spread model are also compared to the forecasts of commercial macroeconometric models.

4.5.1 Quarterly Results

The results in table 16 are meant to be compared to the OLS results for the interest rate specification presented in table 8. Consumption growth and value weighted returns available at t are used to forecast real consumption growth from t to $t + j$. The results suggest that these variables have some explanatory power over all maturities. But from the preliminary data analysis, the differencing of the average data should induce some explanatory power in the consumption autoregression. In the one quarter results, an unadjusted R^2 of .062 is expected. This can be considered a benchmark amount of explanatory power. None of the regressions reported in table 16 can explain much more than the benchmark. The relative forecast power of the interest rate specification and these alternative forecasts is compared in table 18. With minor exceptions, the addition of the fitted values from the interest rate regressions do not provide explanatory power in addition to what is expected as a result of the averaging of the data.

The alternative predictors are also estimated to match the yield spread specification, i.e., to forecast one-step ahead real consumption growth. The differencing of the average data is not a factor in this case.[18] Table 17 provides estimates of these regressions. As one might expect, the ability of these variables to explain growth is reduced. Lagged consumption growth can explain 1–7% of the variation in consumption growth in the overall sample. The stock returns variable is only important in the three quarter forecasts where it accounts for 4% of the variation.

The forecasts from the yield spread model and these alternative predictors are presented in table 19. This table presents 36 comparisons. In 26 of the cases, the interest rate variables have more power than stock returns or the consumption autoregression. In the second sub-period (1972:1–1985:3), the interest rates always win the competition. In the overall sample, it is a close contest between lagged consumption and the interest rate variables. Stock returns are rarely an important explanatory variable.

[18] In the consumption autoregressions, only the one quarter rate of change in consumption is used.

4.5.2 Annual Results

The OLS results for the alternative predictors of consumption growth are presented in table 20. Lagged consumption is particularly important in the overall period (1901–1984) but the explanatory power deteriorates through time. This is due to the high amount of explanatory power in the 1900–1929 period (not reported) when the data were collected in five year intervals and trended. In the later sub-periods, the explanatory power does not exceed what is expected from the averaged data. Lagged stock explain about 1% of the variation in consumption growth in the post-depression sample.

A comparison of the explanatory power of the two alternative predictors predictors of consumption growth and the interest rate specification is presented in table 21. In the full sample, lagged consumption dominates because of the trending of the early consumption data but in the post-1934 samples, the nominal based predictions are more important. As with the quarterly results, the stock returns are rarely important.

Table 22 compares the power of the alternative variables to the predictions of the yield spread specification. The consumption autoregression again wins in the overall period as a result of the interpolated data. The yield spread specifications do better in the post-1934 samples where over 10% of the variation in consumption growth can be explained.

4.5.3 Out-of-Sample Evaluation

The analysis to this point has concentrated on within-sample forecast evaluation. Table 23 presents some out-of-sample forecast evaluation results for the quarterly data. Four models are tested: the real yield spread model, lagged consumption, real value weighted returns and the nominal yield spread model. The evaluation period is 1976:1–1985:1. This period was chosen because it coincides with the period used by McNees (1985) to evaluate the commercial econometric forecasting services. Each model is estimated using ten years of data. The parameters are re-estimated in each quarter of the evaluation period and one to three step ahead forecasts are made. Table 23 presents the RMSE and MAE of these forecasts. The first panel shows the one quarter ahead forecast evaluation statistics. Both the lagged consumption model and the real returns model do slightly better than the yield spread model. In the two and three quarter ahead forecasts, the yield spread model out-performs both the lagged consumption model and the real value weighted stocks model.

Figure 13 provides plots of the one to three step ahead forecasts. The top panel shows the one quarter ahead forecasts. Both the lagged consumption and the value weighted returns forecasts are just the mean of the series. The yield spread model predictions pick up some of the peaks and troughs. The middle and lower panels show the two and three step ahead forecasts. The yield spread predictions seem to move much closer to the actual series. All of the models miss the severity of the 1980 recession.

The analysis of the out-of-sample forecasting performance with the annual data is also contained in table 23. Both of the alternative forecasting variables out-perform the yield spread in the nine year evaluation period. A plot of the predicted values is shown in figure 14. As with the quarterly data, the alternative predictors pick up the mean of the series. This is expected if the slope coefficient in the regression is zero. In contrast, the yield spread series moves with the actual real consumption growth picking up some of the peaks and troughs.

In summary, the out-of-sample forecasting evaluation, suggests that there is some power in the yield spread specification to predict consumption growth – particularly in the two and three quarter specifications.

4.5.4 Comparison to Commercial Macroeconomic Forecasting Services

Consumption of non-durables and services is just a proxy for the true consumption that enters the representative agent's utility function. It is difficult to evaluate the quality of this proxy. Indeed, other macroeconomic variables may be better proxies for true consumption. This section measures the ability of the yield spread based model to predict other macroeconomic variables. The predictions are then compared to the forecasts of the commercial econometric models.

The yield spread model is applied to four macroeconomic variables that are candidate proxies for true consumption. The four variables selected were: real growth in gross national product, nominal[19] growth in total consumption, change in housing starts and growth in real non-residential fixed investment. The growth in these measures can be considered a proxy for the marginal rate of substitution or for the growth in portfolio of productive assets. Both Cox, Ingersoll and Ross (1985b) and Breeden (1986) have linked the term structure to production in the economy. Their models suggest that the real interest rate is positively related to production growth.[20] The growth in GNP and the change in non-residential fixed investment may more appropriately be considered proxies for production growth.

Predictions for these variables are generated in the 1976:1–1985:1 period by re-estimating the parameters at each point in the time series. This presumably is what the large econometric models do either formally of informally. There are a number of problems in comparing the forecasted results with the large econometric

---

[19] The corresponding real measure was not available.

[20] Holding a number of factors constant, such as the variance of the production process, and assuming (among other things) that constant returns to scale production technology, Cox, Ingersoll and Ross (1985 equation 14) and Breeden (1986 equation 15) link production to the term structure.

models. The first is a timing problem. The yield spread model uses information at the beginning of time t to forecast growth from $t+1$ to $t+j$. Although the "early" quarter forecasts are used, the econometric models have more recent information. A second factor is revisions in the data. The yield spread model is fit with revised data while the econometric models do not have these data available. The forecasts of Data Resources Inc. (DRI) and Chase Econometrics are used for comparison. The forecast evaluation covered the 1976:1 to 1985:1 period.

Table 24 presents the summary statistics for the predictions. All of the growth measures are annualized and in percentage terms. The yield spread model dominates the forecasts of the two competing econometric models in terms of growth in total personal consumption. It is probably safe to assume that the large scale models have inflation forecasting power that matches the univariate time series model presented in the paper. If this is true, then it appears that the model would also outperform DRI and Chase's real consumption forecasts. The large econometric models do better than the yield spread model in forecasting most of the other variables – but the difference is usually quite small. The mean absolute errors of the Chase Econometrics model are .2 to .4% better than the two variable yield spread model for growth in real GNP. The RMSE of the yield spread model are lower than the Chase model in all forecast horizons for housing starts. The DRI forecasts are generally better than the Chase forecasts. The yield spread model fares poorly against the other models in the prediction of non-residential fixed investment.

In summary, the model suggested by the consumption-based theory exhibits some ability to forecast growth in the economy. It generally dominates other candidate predictor variables such as lagged consumption growth and real stock returns in terms of within-sample and out-of-sample evaluations. The spread also seems to have some power to predict other procyclical macroeconomic variables. The explanatory power seems non-trivial considering how well it fares against the commercial econometric models.