

Section 7: Results of Model Simulations

Regardless of the mathematical sophistication of the variables incorporated in a model, the accuracy of the calibration process used to determine the parameters, and the timeliness of the values on which the calibration is based, the most important test of the validity of any model is the reasonability of the results. This section will examine the results of a representative run of the model and compare these results with historical values for interest rates, inflation, stock and real estate returns, dividend yields and unemployment rates.

The model was run to generate 5,000 iterations using the base parameters and with the box labeled “Do Not Allow Negative Nominal Interest Rates” checked. The results are presented in several different ways. Table 1 lists the mean levels for the first (one month or one year) and last (50 year) values for key variables, and the 1st and 99th percentile levels for a mid time period (10 years). The figures show “funnel of doubt” plots for different variables that indicate the level of uncertainty that surrounds the output. These “funnel of doubt” graphs (a very descriptive term coined by Redington in 1952) are referred to as “summary graphs” in @RISK. These graphs can be generated in @RISK for any array of a single variable over time. The x axis indicates the time period, with larger values representing times further in the future. The y axis indicates the value assumed by the variable. The “funnel of doubt” graphs show the mean value for the 5,000 iterations (solid line) the 25th and 75th percentile values (dark shaded section) and the 1st and 99th percentile values (lighter shaded section). Expanding funnels indicate that the values become more uncertain the longer the time horizon over which the projection occurs. Narrowing funnels indicate that the variables become more predictable when forecasting the long-term movements of a particular variable.

Next, histograms are used to illustrate the full probability distribution of the values for a particular variable at one point in time. In addition, the distribution of historical values, where appropriate, is plotted on the histograms to show how the model results compare with the actual values for that variable. Finally, correlation matrices are shown for the simulation values and for historical values.

Table 1 lists the results for selected key variables. These include the Real and Nominal Interest Rates, the Inflation Rate, the Unemployment Rate, Large Stock Returns, Small Stock Returns, Dividend Yields and Real Estate Returns. For Interest Rates and Inflation, annualized values for 1 Month, 1 Year and 10 Year periods are shown. The first column indicates the first value the model generates for these variables. For both Interest Rates and Inflation, this is the value assumed after one month. For the other variables, the results are shown after the first year. (Since the stock and real estate returns are annualized, the 1st and 99th percentile values of the initial monthly returns for these variables represent unrealistic annual values that are not really indicative of the results the model would predict over a full year.)

Real Interest Rates

The results for Real Interest Rates indicate that the mean value for the first month is 0 percent. By the end of the 50 year period modeled, this value has moved to 3.0 percent. This result is

entirely in line with the specifications of the model. The one month value would be closely aligned with the initial short-term interest rate (r_{init1}), which is 0% in the base parameters. The value would just begin to revert to the long-term mean after one month. The mean of the final value in the results, after 50 years, is around the mean reversion level for the long rate ($rm2$), which is 2.8 percent.

To provide an idea about the range of values this variable assumes, the 1st and 99th percentile values of the 1 month real interest rate after 10 years are displayed in columns 3 and 4. In 1 percent of the iterations, the 1 month real interest rate, on an annualized basis, is negative 5.3 percent. On first observation, this result seems nonsensical. Why would an investor be willing to lose money, in real terms, by investing at a negative real interest rate? Instead, an investor would just hold cash rather than lose 5.3% per year, after adjusting for inflation. However, this may not be as unrealistic as it seems. First, this is only the 1 month rate, and the annualized rate over the month represents a negative 0.4%. Second, this return may represent the best return available. If inflation is high, then holding cash would generate an even larger loss. In times of high inflation, the best real return an investor can receive may be negative. Finally, real interest rates are not observable. The true real interest rate is the return required, over and above expected inflation, for the specific interval. However, no one knows the expected inflation rate, as determined by the financial market, at any point in time.

In practice, two approaches have been used for estimating the expected inflation rate. In one, economists' forecasts of inflation have been used to measure the expected inflation rate. Economists, though, do not represent investors. By training and occupation, the economists included in the surveys are not at all representative of the general financial market participants. Investors may consider some economists' forecasts in making their own determination of what to expect regarding future economic conditions, but many other factors, including their own experience, the counsel of other participants, and recent historical experience, are used to determine their inflation expectations. There is no survey of representative market participants to determine what they truly anticipate for the inflation rate.

The second approach has been to examine actual inflation rates that have occurred, and then subtract those from prior interest rates. This approach is termed an ex-post analysis, since it examines what actually occurred over a particular time interval. This approach is also flawed for several reasons. First, there is no reason to believe that the market is prescient regarding inflation expectations. Especially in the case of an unexpected shock to the system, such as oil price increases during the 1970s, the market does not know what will happen in the future. It cannot even be assumed that errors in forecasting will cancel out over time, since the market could be biased to under, or over, estimate future inflation. Second, actual inflation cannot be measured. The Consumer Price Index and other values commonly used to determine inflation are widely recognized as being inaccurate. These indices track the prices of specific goods and services that are not completely representative of the entire economy. Also, these indices cannot recognize the substitution effect that consumers engage in continually, such as buying more chicken than usual when beef prices rise more rapidly than chicken prices, or driving less when gasoline prices soar. Due to these problems, it is not possible to claim that real interest rates cannot be negative, so a small negative value over a short time interval does not necessarily represent a problem.

On the opposite side, the 99th percentile value for 1 month real interest rates after 10 years is 10%. The same limitations described above also apply to this value.

For 1 year real interest rates, the mean of the first value, one month into the future, is 0.3%. This reflects regression from the initial value of 0% to the long-term mean of 2.8%. The mean of the last value, after 50 years, is, in line with these parameters, 2.9%. The 1st – 99th percentile range after 10 years is a negative 5.1% to 9.7%, reflecting a similar distribution for the full year as was observed for the monthly values. For the 10 year real interest rates, the means of both the first value is 1.1% and the final value is 2.6%, reflecting the strength of the mean reversion over this long a period of time. The 1st – 99th percentile range after 10 years is negative 3.3% to 7.6%, reflecting the more compact distribution for long-term (10 year) real interest rates, compared to shorter time horizons.

The Funnel of Doubt graphs of 1 month, 1 year and 10 year real interest rates, Figures 1-3, all reflect the same shape, although the scaling differs. These graphs do indeed resemble a funnel with a narrow neck. This pattern occurs because the first 12 points represent monthly intervals, which have small changes in values, and the latter steps are annual or larger intervals, which lead to correspondingly larger changes. The level of uncertainty increases over the entire 50 year time frame, but the shifts toward the end of the simulation period are less pronounced. This shape occurs due to the structure and parameterization of the model. The stochastic term generates the initial spread of the distribution, but the mean reversion impact offsets this tendency, keeping the “funnel of doubt” from expanding further extensively.

Inflation

The next variable of interest is the inflation rate, which is added to the real interest rate to determine the nominal, or observed, interest rate. As shown in Table 1, the mean value of the 1 month inflation rate is 1.1% after only one month, and 4.8% after 50 years. This is in line with the base parameters of 1.0% for the initial inflation rate (qint1) and 4.8% for the long-term mean (qm2). The mean reversion speed for inflation is much lower than the mean reversion speed for real interest rates (0.4 compared to 1.0). After 50 years the monthly inflation rate averages out to be the long-term mean. The 1st – 99th percentile range after 10 years is -5.3% to 14.5%. Inflation can be negative (which would be termed deflation) and small negative monthly values have occurred in recent years. Monthly inflation values in excess of 14.5% did occur during the 1970s.

The mean 1 year inflation rate begins at 1.6% and moves to 4.8% by the end of 50 years, again both in line with the model parameters. The 1st – 99th percentile range of the 1 year inflation rate after 10 years is -3.7% to 12.9%. Although the United States has not experienced deflation over an entire year since 1954, any realistic model must assign a significant probability to this occurring again.

The mean 10 year inflation rate begins at 3.6% and moves to 4.5% by the end of 50 years. These are closer to the long-term mean parameter of 4.8%, but fall below due to the issue of geometric means discussed in the Users Guide. The 1st – 99th percentile range of the 10 year inflation rate

after 10 years is 2.0% to 6.9%, demonstrating that, for longer time horizons, the geometric average of inflation is less variable.

The Funnel of Doubt graphs of 1 month, 1 year and 10 year inflation rates are shown on Figures 4-6. The uncertainty of the 10 year inflation rate is much smaller than it is for 1 month and 1 year, reflecting the strength of the mean reversion term for this single factor model. Although inflation varies widely over shorter time horizons, in this model the long-term inflation rate is much less variable. This pattern can be altered by increasing the volatility of the inflation process (qs_1) or reducing the mean reversion speed (qk_1).

The histograms for the 1 year projected inflation rates and of actual 1 year inflation rates from 1913-2003 (from January to January) are shown on Figure 7. It is readily apparent that the modeled inflation rates generate a nice bell shaped curve, whereas the actual inflation rates are much less smooth. One reason for this difference is that the model results are based on 5,000 iterations, while the actual data contain only 90 data points. More importantly, though, the projected values are derived from a concise mathematical expression that will produce a smooth distribution of results, but the actual inflation rates depend on the interactions of an almost unlimited number of variables. The key question, though, is whether the model adequately expresses the probability distribution of potential inflation rates. The actual inflation rates are more leptokurtic (fatter in the tails than a normal distribution) than the modeled values, but reflect the central portion of the graph fairly well. All of the large negative inflation rates occurred prior to 1950. Many of the positive outliers are from years prior to 1980, when monetary policy was less focused on controlling inflation. If the user wishes to obtain a distribution for inflation that reflects the long-term experience, then the inflation parameters should be revised to generate a more diffuse distribution.

Nominal Interest Rates

Nominal interest rates are the sum of the real interest rate and the inflation rate. The mean values for 1 month nominal interest rates were 1.1% for the first month and 7.8% for the 50th year. The initial nominal interest rate indicated in the model (1.1%) is in line with the current (June, 2004) level of 1.1%. Since inflation is running in the range of 2.5% (2.3% for the entire year 2003, and at the 2.1% level for the first five months of 2004), the current nominal interest rate reflects either an expected reduction in inflation (or even deflation), a low (or negative) real interest rate, or some combination of these two effects. The 1st – 99th percentile range for the 1 month nominal interest rate after 10 years is 0% to 19.4%. In this simulation, nominal interest rates were restricted to be not negative to generate results in line with economic values.

There has been significant debate over the proper way to deal with negative nominal interest rates in interest rate models. Some modelers have set boundary conditions that prevent nominal interest rates from becoming negative. This is one of the options available in running this model. Other modelers have not been concerned over negative interest rates, either because the mathematical characteristics of the model are more important than the practical applications, or the incidence of negative nominal interest rates is too infrequent to require significant attention.

The occurrence of negative nominal interest rates can be problematical, since the purpose of this research is to develop values that can be used in applications for insurance companies and others. However, artificial constraints on the values do reduce the connection of this model to the research underlying the structure of the interest rate model. The model allows users to set a minimum value of zero for the nominal interest rate, or to accept negative nominal interest rates, depending on the application.

The mean 1 year nominal interest begins at 1.9% and moves to 7.7% by the end of 50 years. The initial value is again in line with the current level of interest rates. The 1st – 99th percentile range of the 1 year nominal interest rate after 10 years is 0% to 18.3%.

The mean 10 year nominal interest begins at 4.6% and moves to 7.1% by the end of 50 years. The initial value is in line with the current level (July 2004) of interest rates for long-term bonds, as 10 year U. S. Treasury bonds have a yield to maturity of 4.4%. The 1st – 99th percentile range of the 1 year nominal interest rate after 10 years is 0.6% to 12.7%.

The Funnel of Doubt graphs of 1 month, 1 year and 10 year nominal interest rates, Figures 8-10, are similar to the real interest rate and inflation graphs, but have a barrier at zero since the restriction that nominal interest rates not be negative is applied in this case. This restriction is illustrated by the 1st percentile line on Figures 8 and 9, but not for the 25th percentile line. The effect of the restriction is not apparent for the 10 year nominal interest rates. The level of uncertainty increases over the 50 year time period used in the forecast. The further into the future the forecast of nominal interest rates is made, the more uncertainty surrounds the estimates. Since the nominal interest rate is determined by adding the real interest rate to the inflation rate, the increasing uncertainty reflected by real interest rates and the inflation rate generates the same behavior for nominal interest rates.

The histograms for the 3 month, 1 year and 10 year model nominal interest rates and the actual 3 month, 1 year and 10 year nominal interest rates are displayed in Figures 11-13. (The 1 month values are not available for historical data consistently over a long enough time period to be relevant. The model can generate 3 month interest rates, which were used for this graph.) The modeled interest rates are influenced by the initial parameter values, particularly the initial real interest rate (rinit1), the initial mean reversion level for the real interest rate (rinit2), and the initial inflation level (qinit1). The influence of these initial values diminishes over time, so histograms of these values selected further into the future would differ from those based on more contemporaneous values. Based on the base parameters selected, the values will increase over time. In this case, the model distributions in Figures 11-13 are based on values after one year.

Significant differences do exist between the modeled and historical distributions for interest rates. In Figure 11, the modeled 3 month nominal interest rates are 0 in almost 20% of the cases, whereas actual 3 month interest rates have never been below 0.5 percent (the column reflecting the 1% bin represents values between 0.5 and 1.5 percent). However, combining the model values for 0 and 1 percent indicates a total very much in line with actual values. In addition, the model distributions are smoother than the actual values, which is natural since the model results are based on 5,000 iterations whereas the actual results, even though derived from 845 (monthly) or 614 (1 and 10 year) observations, are not at nearly as smooth, indicating that the system that

generates interest rates is not as straightforward as the model. Modeled interest rates are generally lower than the historical rates, which results from starting the model based on the current level of interest rates, which are lower than historical averages.

For the 1 year interest rates, shown in Figure 12, the modeled rates are notably lower than historical rates. This difference is due to setting the initial values at the current level of interest rates, which are much lower than historical average values. Since the model starts with current interest rates, it projects interest rates that, even one year from now, tend to be lower than historical values.

The comparison between the 10 year modeled rates and the 10 year historical rates, Figure 13, indicates a few differences. The modeled interest rates are more compact than actual 10 year interest rates have been. If the user feels that the variance of the model values should be closer to the historical distribution, then the strength of the mean reversion factor in the interest rate model can be reduced, but this would increase the incidence of negative interest rates unless the user selects to avoid negative nominal interest rates. The other significant difference is the skewness. The historical rates exhibit positive skewness, but the modeled rates have a slight negative skewness. Finally, the model rates are lower than historical values, again due to starting with the current low levels of interest rates.

Stock Returns and Dividends

The values for large and small stock returns indicate, as expected, higher average returns and greater variability for small stocks. As shown on Table 1, the mean of the initial values (after one year) of large stocks is 8.7% and of small stocks is 13.4%. The mean of the large stock values increases to 11.6% at the end of 50 years and for small stocks increases to 13.6%. The 1st – 99th percentile range after 10 years is –15.9% to 29.6% for large stocks and –15.9% to 39.7% for small stocks.

The Funnel of Doubt graphs, Figures 14 and 15, indicate an inverted funnel, compared to the displays of interest rates and inflation. This means that uncertainty reduces over time and is due to the way the values are calculated. The values listed for large and small stock returns are those projected to be earned over the relevant time periods, shown in annual terms. The 1 year values are returns over a one year period. The 10 year values are those earned over a ten year period, etc. Since the individual draws are independent, the variance reduces over longer time periods. Thus, the average annual returns expected over a 50 year period are much more predictable than those for a 1 year period.

Stock returns are determined based on a random draw from one of two distributions. For both large and small stocks, two regimes exist. For large stocks, the mean values in each regime are 0.8% monthly return for regime 1 and -1.1% monthly return for regime 2 with volatilities of 3.9% for regime 1 and 11.3% for regime 2. For small stocks the mean values are 1.0% and 0.3% and the volatilities are 5.2% and 16.6% for the two regimes. Which regime applies during any month is determined randomly, based on indicated switching probabilities. The regimes switches are correlated, so if large stocks are in the low volatility regime, then small stocks are more likely to be in the low volatility regime as well.

Histograms of the 1 year returns for the large (Figure 16) and small (Figure 17) stock returns as generated by the model are displayed, along with actual 1 year returns for 500 large stocks for 1871-2004 and small stock returns over the period 1926-2003. The large stocks are based on the S&P 500 (or a sample chosen to behave similarly for the years prior to the construction of the S&P 500). The data are available online at a website generated by Robert Shiller, author of *Irrational Exuberance* (http://www.econ.yale.edu/~shiller/data/ie_data.htm). The small stock values are based on Ibbotson's *Stocks, Bonds and Bills*. The graphs for large stocks (Figure 16) are relatively similar, although, as would be expected, the results of the 5,000 iterations of the model produce a smoother distribution. The histograms for small stocks (Figure 17) show that historical values have been more variable, with a notable outlier at 190% return, which represents a single observation. The model values also have single observations around that level, but no one bin produces as large a proportion of the outcomes as the one occurrence out of 78 years of the historical experience to be as obvious on the graph.

The dividend yield for equities is 1.5% for the first year and 2.3% for the last year values. The 1st – 99th percentile range after 10 years is 0.6% to 3.9%. The Funnel of Doubt graph of the dividend yield, Figure 18, increases over time as interest rates and inflation do. Figure 19 displays the histogram of the modeled dividend yields and the actual dividend yields over the period 1871-2003, based on data available from Robert Shiller. Historically, dividend yields have varied more widely than the model predicts and have been centered at a higher level. This is the result, in part, of a long-term decline in dividend yields. According to Peter Bernstein (*Against the Gods*, p. 183) prior to the late 1950s dividends tended to be higher than interest rates on corporate bonds. This was based on the understanding that stocks were riskier than bonds and therefore should pay a higher return. Since 1959, though, dividend yields have tended to be lower than interest rates, ranging from 1.1% to 5.4%, which is more in line with the model.

Unemployment and Real Estate Returns

The mean value of the unemployment rate, as shown on Table 1, begins at 6.0% and increases to 6.1% (which is the long-run mean value) for the end of 50 years. The 1st – 99th percentile range after 10 years is 3.5% to 8.7%. Figure 20 shows the Funnel of Doubt graph, neither increases over time (as interest rates and inflation do) nor decreases (as stock returns do). The unemployment rate is modeled as a one factor mean reverting value, with the initial unemployment rate (u_{init}) set at 5.0%, the mean reversion level (u_m) set at 6.1% and the monthly mean reversion speed (u_k) set at .132 in the base parameters. By the end of 1 year (the starting point for this graph), the mean value has attained the 6.0% level and the spread has reached the long-run level as well, resulting in a stable funnel. The unemployment rate also has a negative correlation with changes in inflation. The histogram of modeled unemployment rates, along with the distribution of actual values over the period 1948-2003 are shown in Figure 21. The historical values represent the unemployment rate each January from 1948-2004. By selecting only a single unemployment rate from each year, the frequency of the historical values corresponds with that of the model values, which are the unemployment rates indicated after the first year of the model run. Although the actual unemployment rates have varied a bit more than the model results do, the distributions are relatively close.

Real estate returns are the final variable included in this model. From Table 1, the mean value of real estate returns is 8.1% in the first year and 9.4% after 50 years. The 1st – 99th percentile range after 10 years is 3.0% to 16.1%. The Funnel of Doubt graph, Figure 22, is similar to the returns on stocks, for the same reasons. The histograms of modeled results and the actual returns based on the National Index from National Council of Real Estate Investment Fiduciaries (NCREIF) for 1978-2003 are shown on Figure 23. The model values show a smooth distribution that is centered about the historical returns. Unfortunately, only 26 years of annual returns are available, so it is difficult to draw any conclusions on the fit.

Correlations

Table 2 displays the correlation matrix for all the output variables discussed above 1 year into the simulation (row 16). Table 3 displays the correlations matrix for large and small stock returns, 3 month, 1 year and 10 year treasury interest rates and inflation over the period April 1953 – December 2001, based on the Ibbotson data (for stocks) and St. Louis Federal Reserve Data for interest rates and inflation. Several conclusions can be drawn about the validity of the model based on a comparison of the two correlation matrices. First, the historical correlation between large and small stocks is .744. The correlation between the model values of large and small stocks is .698, which looks quite reasonable.

The correlation between inflation and T-bills has been .593 historically. This correlation is also clearly reflected in the model values, with a correlation of .906 between the one month inflation rate and the 1 month nominal interest rate, .892 between the 1 year inflation rate and the 1 year nominal interest rate, and .617 between the 10 year inflation rate and the 10 year nominal interest rate. Since the nominal interest is the sum of the real interest rate and the inflation rate, and the real interest rate is constrained to be no less than the negative of the inflation rate, this correlation is built into the model.

Historically, T-bill rates and stock returns have been negatively correlated (-.078 for large stocks and -.065 for small stocks). In the model, there was a slight positive correlation between the 1 year nominal interest rate and stock returns (.099 for large stocks and .087 for small stocks). Also, the historical correlation between inflation and stock returns has been negative (-.138 for large stocks and -.100 for small stocks). The correlation in the model values between the 1 year inflation rate and large stocks was .089 and .076 for small stocks. Although the signs of the correlations differ, these small values are not particularly significant.

Summary

This model provides an integrated framework for sampling future financial scenarios which represent a reasonable approximation to historical values. The model should prove useful for a variety of applications, including dynamic financial analysis, dynamic financial condition analysis, pricing embedded options in policies, solvency testing and operational planning. The model produces output values for interest rates, inflation, stock and real estate returns, dividends and unemployment. However, there are several areas that may be a cause of concern to some users. The model inflation rates have fewer outliers than long-term historical experience. Users

should decide if the recent inflation experience is considered to be the best predictor of future inflation, or if the parameters should be changed to reflect these expectations.

During the Depression, unemployment rates in the United States were around the 25 percent level. Since 1948, rates have been much lower, and the model values correspond fairly closely with this recent experience. However, users of the model may want to establish parameters for the unemployment model that reflect earlier, and more volatile, experience.

The base parameters provide one feasible set of values to use in modeling future economic conditions. These should be viewed as a starting point in these applications. However, users should develop an understanding of the impact of the different parameters and then adjust these parameters as necessary to generate distributions that are suitable for the particular applications of the model.