PROJECT TEMPLATE: MUNICIPAL BOND SPREADS

Municipal bond yields give data for excellent student projects, because federal tax changes in 1980, 1982, 1984, and 1986 affected the yields. This project template for municipal bond spreads suggests ideas for a student project. We focus on the statistical analysis. You do not need to know corporate tax law to complete this student project.

Step #1: Define the Spread

You can define the municipal bond spread two ways for the student project:

~ Municipal bond yield minus Treasury bond yield (for bonds of similar duration): This spread is negative, since the tax advantage of municipal bonds outweighs their higher risk.
~ Municipal bond yield minus corporate bond yield for bonds of equivalent grade: This spread is more negative, because corporate bonds have higher yields than Treasury bonds.

You can use either definition. The first definition is more common; the second definition adjusts for economic conditions that widen or narrow the spread.

If you prefer to work with positive figures, use the Treasury bond yield or the corporate bond yield minus the municipal bond yield.

Take heed: Examine if a multiplicative spread is more stable than an additive spread. Either type of spread is fine for the student project.

The spread reflects tax law, which is a multiplicative factor applied to the pre-tax yield. The tax advantage is twice as great when the pre-tax yield is 12% as when it is 6%. You may use multiplicative spreads or spreads of the logarithms of the yields. No method is perfect.

We do not tell you the optimal way of forming the spread to get a stationary time series. You may try two methods and compare their correlograms.

Financial economists do not agree on the best definition of the spread. You might use a simple additive spread or a multiplicative spread. Ideally, use a method that gives a stationary time series.

Step #2: Graph the Spread

Using a spread offsets the distorting effects of changing inflation rates and interest rates. The changing trends, means, and volatilities of interest rates don’t appear in the spreads. This makes the modeling easier. An AR(1), AR(2), MA(1), or ARMA(1,1) process should fit well.

Step #3: Spread Eras

Tax law relating to municipal bond yields changed in 1980, 1982, 1984, and 1986. A time series spanning all years is not stationary, since the spread narrowed between 1980 and 1986.

Your student project can examine the graph of the full series and fit an ARIMA process to one time period. You can also compare the ARIMA process for pre-1980 vs post-1986.

~ In 1986, the corporate tax rate declined from 46% to 35%, reducing the tax advantage of municipal bonds, and raising the required yield on municipal bonds.
~ Tax law changes in 1980, 1982, and 1984 eliminated almost all the tax advantages for commercial banks. Banks bought over 50% of municipal bonds before 1980, but they buy about 1% now. The elimination of the tax advantage for commercial banks raises the required yield on municipal bonds.
~ Tax law changes in 1986 reduced by 15% the tax advantage for property-casualty insurance companies, who had been the second largest clientele for these bonds before 1980 and the largest clientele after 1984. They remain the primary clientele even now, since other investors have less use for these bonds. The reduction in their tax advantage raises the required yield on municipal bonds.

Step #4: Fit ARIMA Processes

Fit an ARIMA process to (i) the years before 1980 and/or (ii) the years after 1986.

Use the NEAS step-by-step guide on the discussion board. Some items are simpler for municipal bond spreads.

Seasonality has almost no effect, for three reasons:

~ The FED adjusts the money supply to remove the seasonality in short rates.
~ Long duration rates smooth any remaining seasonality.
~ The spread eliminates any seasonality that survives the smoothing.

The trend of interest rates is much reduced in the spreads and may not appear at all. The reduction (or elimination) of the trend differs for additive vs multiplicative models. Using an additive model on the logarithms of the rates gives a multiplicative model.

Taxes are multiplicative, but they apply only to the interest component of the rate. If the pre-tax yields of two bonds are 20% and 30%, for a 10% difference, the after-tax yields are 13.0% and 19.5%, for a 6.5% difference. If the 20% bond is tax exempt, the after-tax difference is –0.5%.

For the student project, think through how the municipal bond tax exemption might be modeled as a function of the municipal bond pre-tax yield. Instead of fitting an ARIMA process to the municipal bond spread, a financial economist might fit the process to the difference between the spread and the tax exemption. Do not worry about the optimal method. Explain what you do and show how you use the statistical tools.

Financial economists do not agree on the best way to define the spread or model the spread. For the student project, you can use a simple additive model or a complex multiplicative model using differences between the spread and the tax exemption.

The optimal ARIMA process depends on your definition of the spread. A simple additive model shows a discrete jump in 1986 to 1987. If you use monthly rates, the jump occurs on August 6, 1986. Examine your graph of the spread, identify the eras, choose the time period or periods you will model, and fit ARIMA processes.

Step #5: Compare the Processes

The pre-1980 and post-1986 eras have different means. You can compare the means and relate them to the tax advantages in the two eras. The tax law is intricate, and the student project is not graded on your analysis of tax law. We examine if you properly apply the statistical techniques to fit an ARIMA process to the spreads.

PROJECT TEMPLATE ON HOUSING SALES.

Home construction and sales is highly cyclical, even more than auto sales. A student project may use the following steps to fit an ARIMA process.

Convert housing sales to real dollars (divide by CPI). If the housing sales are in units, not in dollars, you need not convert to real dollars. The Excel work-sheet has sales in units.

Regress housing sales on real per capita personal income. If you use housing sales by state, use state personal income; otherwise use countrywide personal income. The Excel work-sheet has countrywide housing sales.

You may find a leveraged effect: A 3% increase in real per capita personal income causes a 6% rise in new home sales. Fit an ARIMA process to the residuals of this regression.

We speak of built-up demand for new homes. You may compare retail sales of several industries (homes, autos, electronics, furniture, food). Compare the effect of personal income on each industry and the optimal ARIMA model for each industry.

Real per capita personal income is a good explanatory variable, but it is not the only one. You may also use GDP and unemployment rates. But don’t use these in combination, or the explanatory variables will be correlated and the regression will not work well.

Housing sales depend on mortgage rates. You may use a multiple linear regression, with real per capita personal income and real mortgage rates as the explanatory variables. You can get mortgage rates from internet sites, or you can use the long-term bond rates as a proxy for mortgage rates.

Housing sales are seasonal. De-seasonalize the data or use a 12 month autoregressive parameter. See the project template on auto sales. Many of the comments about retail car sales apply to home sales as well.

Home sales are sub-divided into “not started,” “under construction,” and “completed.” You can compare the fitted ARIMA process for each series. We expect that sales of homes not yet started depend more on fluctuations in personal income. You can test this hypothesis.

Project Template on GNP and GDP

GNP and GDP are measures of economic activity.

- GDP, or Gross Domestic Product, is economic activity by persons working in the country, including foreigners. It does not include citizens working abroad or the income from foreign capital. GDP is the measure now used for a country's economic activity.
- GNP, or Gross National Product, is economic activity by the country's citizens. It includes income by citizens from working abroad (or capital invested abroad) but not income by foreigners working in the country.

GNP and GDP are measured on a real or a nominal level. The raw indices are seasonal, and the index is seasonally adjusted for use with other indices.

Take heed: The indices on the discussion forum are quarterly, seasonally adjusted, in nominal dollars.

You can do a student project on a GNP or GDP index. You may also use GDP or GDP to form structural models.

Illustration: Your student project fits an ARIMA process to new car sales or new home sales. Sales of durable goods, such as autos and homes, vary with economic activity.

- In prosperous years, when the economy is strong, consumers use their savings (or take loans) to buy durable goods.
- In recessions, sales of durable goods decline.

New car sales and new home sales vary from year to year, and no ARIMA process fits well. For your student project, regress auto sales or home sales on GDP or GNP and fit the ARIMA process to the residuals of the regression.

Many of the project templates discuss structural models by regressing the time series on GDP (and other macroeconomic indices). If the time series you are examining is monthly, interpolate the quarterly GDP series to a monthly series.

Illustration: Suppose the first quarter 20X7 GDP is $210 trillion and the second quarter 20X7 GDP is $222 trillion. Use these figures for February and May 20X7. Interpolate for March 20X7 as $214 trillion and for April 20X7 as $218 trillion.

Time Series Project Template: LIBOR Rates

LIBOR, or "London Inter-Bank Offered Rate." is the rate that large London banks offer each other for inter-bank deposits. Deposits are funds loaned to banks, so LIBOR is the rate at which a fellow London bank can borrow money from other banks.

- LIBOR rates incorporate variables such as time, maturity and currency rates.
- Hundreds of LIBOR rates are reported each month in numerous currencies.

LIBOR is a risk-free rate with maturities of twenty-four hours (overnight) to five years. It is reported each day at 11 a.m. London time and then fluctuates based upon the market's expectations for economic activity and the future direction of interest rates.

LIBOR loans are expressed in Eurodollars: United States currency held by foreign entities, such as a British or German banks or insurance companies.

Eurodollars reflect American firms paying dollars internationally-domiciled firms for goods, service, and merchandise purchased.

Choose a LIBOR Rate

The NEAS web site has 7 LIBOR for U.S. dollars, ranging from over-night to one year.

- LIBOR rates with maturities of one month or longer have the longest history (21+ years) and less day-to-day fluctuation. If you want a smooth time series with many observations, choose six month or twelve month LIBOR.

- LIBOR rates with short maturities are mostly likely to show seasonality and cycles, and they reflect other macroeconomic variables. To examine the effects of U.S. GDP or the demand for money on interest rates, use over-night LIBOR.

Take heed: We do not say that over-night LIBOR definitely shows cycles or seasonality. Examine the ARIMA process of over-night rates, and explain what you conclude.

The actuarial examinations cover financial economics: forward rates at various maturities. The Eurodollar market is extremely liquid and closed tied to the LIBOR market. For a student project, you can search the internet for Eurodollar futures rates and compare those rates to LIBOR rates.

LIBOR rates are expressed in numerous currencies. We show only U.S. dollar rates, but you can search the internet for other currencies as well. You can do a student project on interest rate parity, putting together exchange rates and risk-free rates in each currency.

Day Index and Seasonality

Rates are shown for business days. You may examine two types of seasonality:

- Weekly seasonality: the average Monday rate vs the average Friday rate
- Annual seasonality: using the business days in the year.

The number of business days varies from year to year, ranging from about 242 to 250.

The day of the week is clear from the file. Rows are sets of five consecutive days: Monday through Friday.

Take heed: Excel has built-in functions to identify the day of the week.

- Convert the year, month, and day to an Excel date: =DATE(year, month, day).
- Convert the date to a weekday: =WEEKDAY(date).

Excel also computes the calendar days and the workdays between any two dates.

For annual seasonality, use five day moving averages.

Illustration: January 10 may be a weekday some years and a weekend other years. Instead of January 10, use a five day moving average of January 8 – January 12.

MATURITY SPREAD

You can form a time series of LIBOR maturity spreads, such as 12 month LIBOR minus 1 month LIBOR. The maturity spreads eliminate some of the random fluctuation in LIBOR rates, leaving a more stable time series for ARIMA modeling.

Take heed: You may also use (1 + 1 month LIBOR) / (1 + 12 month LIBOR)

REAL LIBOR RATES

For a time series of real LIBOR rates, use 1 month LIBOR – CPI inflation the previous month. (Alternatively, divide by 1 + the inflation rate.)

The nominal interest rate depends on expected inflation. We do not have a daily index of expected inflation, so we don’t have a daily time series of real LIBOR rates.

PROJECT TEMPLATE ON ARIMA MODELING OF PERSONAL INCOME.

The Excel spread-sheet on personal income shows per-capita income by state over the past century. Each state time series – by itself – is probably a random walk, with a drift of inflation plus a productivity increase.

You might begin a student project by analyzing one state index. Nominal personal income is not stationary. You may take logarithms and first differences to see if the resulting time series can be modeled by an ARIMA process.

Inflation distorts the time series. Converting the nominal dollars to real dollars should leave a time series of productivity growth. Taking logarithms and first differences is now more likely to produce a stationary series.

Personal income, GDP, and unemployment are related. You can form a structural model and fit an ARIMA process to the residuals. You may have to search for GDP and other macroeconomic indices that extend back as far as the personal income index.

The macroeconomics on-line course has extensive analysis of personal income by

- State within the United States.
- Territory or region within European countries.
- Country within the world.

The macroeconomics on-line course assumes absolute convergence for the states of the U.S. For your student project, use the relativity of the state per capita income to countrywide per capital income.

Economists believe this relativity is slowly mean reverting – about 2% per annum.
The macroeconomics on-line course call this absolute convergence.

Illustration: If State W's per capita income is 150% of countrywide in 20X7, it is expected to be 149% in 20X8.

For any one state, the relation is hard to see, since political changes may raise or lower per capita income.

Illustration: An increase in the state income tax or more onerous state regulation (New York, New Jersey, California) may slowly reduce average personal income over the next 40 years, whereas more business-friendly taxes and regulation may increase average personal income (Florida and other Southern states).

You may examine how state fortunes, such as the decline of the auto industry (Michigan in the 1970's) or the high immigration into California and Texas in the past two decades have affected relative personal income.

An ARIMA process in one state does not capture these political effects. A single state may not have a stationary time series. Start your student project with one state. You can fit an ARIMA process to each state, but you probably won't get ϕ_1 = 98%. Instead, you may model all the states. Use the Excel built-in functions for a linear trend, regressing each state's income relativity on the relativity one year back. With 77 years and 48 states in the Excel work-sheet, you have enough data to estimate the ϕ_1 parameter.

Take heed: Estimating the ϕ_1 parameter can be difficult. Instead, your student project can assume ϕ_1 = 98% and test if the residuals from the ARIMA process are white noise.

Each state has 77 time series observations. An AR(1) process leaves out the multi-year effects. You may compare AR(1), AR(2), and ARMA(1,1) processes.

PROJECT TEMPLATE ON ARIMA MODELING OF RETAIL CAR SALES

The textbook for the on-line time series course has many illustrations of ARIMA modeling for industry sales. The NEAS discussion forum has numerous time series of retail sales by industry that can be used for student projects.

Each industry differs. The student project should examine trends, cycles, seasonality, and changes in market demand.

The Excel spread-sheet on retail auto sales has six time series. The six series are related:

- Total auto sales = new car sales + used car sales (#3 = #4 + #5).
- Auto + parts sales = auto sales + parts sales (#1 = #2 + #6).

Take heed: The attributes of the industry help you model sales. Auto manufacturers gave large rebates in some years, raising sales above their expected levels.

- Higher new car sales one year may reduce used car sales that year and new car sales the next year.
- The student project does not require special knowledge of the industry. But combining GDP, car sales, and inflation may help you interpret the results.

Inflation affects dollars sales. To convert sales to real dollars, divide by the CPI. It is easier to fit an ARIMA process to real dollar sales than to nominal dollar sales. If you use nominal dollars, you must take logarithms and first differences to create a stationary time series. But the first differences distort the moving average and autoregressive parameters, and make the ARIMA process harder to interpret.

Sales of durable goods (like autos) vary strongly with GDP. A 2% increase in real GDP may cause a 5% increase in auto sales, and a recession (a drop in real GDP) may cause a large decline in auto sales.

We may state this relation as demand build-up. New car sales show demand build-up. If sales are lower than expected in 20X7, demand builds up and sales may be higher than expected in 20X8.

Demand build-up causes a *moving average* ARIMA process, which does not occur in most industries and in single firm sales. Its effect on auto sales is unclear. Your student project may examine if sales of durable goods have a moving average component.

New car sales are seasonal, with new year models appearing in October. But auto manufacturers offer so many inducements (no cash down financing, rebates, special loan offers) that the seasonality can be hard to model. Graph the monthly new car sales, adjust for inflation and trends, isolate any apparent seasonality, and examine 12 month sample autocorrelations. If the seasonality is material, de-seasonalize the data or use 12 month autoregressive parameters.

New car sales affects used car sales, but the relations are hard to isolate.

An increase in new car sales may cause a decline in used car sales. Convert nominal dollar sales to real dollar sales and then regress used car sales on new car sales. Fit an ARIMA process to the residuals.

A downturn in the economy may cause a decline in all car sales. Regress used car sales on new car sales and real GDP. Fit an ARIMA process to the residuals.

The combination of these items makes car sales an excellent student project. You may use lagged variables, such real GDP one month earlier.