To Appear in *The Handbook of Municipal Bonds, Probus Publishing, Chicago 1994.* pp. 441-450.

# Knowledge-Based Approaches for Evaluating Municipal Bonds

Roy S. Freedman
William P. Stahl

## Introduction

Municipal bond analysis requires a sophisticated mix of quantitative and qualitative techniques for trading and management activities. Even though much of the quantitative component is automated and enriched by a tremendous amount of data, much of the critical buying and selling decisions are based on individualized, qualitative judgments.

The actual use of formal quantitative portfolio models vary from institution to institution. At one extreme, typified by rigorous investment constraints, the model recommendations are followed precisely. In most cases, however, the model recommendations are just — recommendations. The model outputs, based on information available to it, are "guidelines." All decisions relies on the judgment, experience and intuition of the analyst or portfolio manager.

To be successful, a municipal bond portfolio manager must maximize return and minimize risk in a dynamic world. Events that can trigger a portfolio rebalancing include (i) changes in the yield curve; (ii) changes in trading relationships, such as beta and duration (iii) changes in credit quality. Of these three events, events involving credit quality are "less quantitative" and involve subjective knowledge. The first two are "more quantitative" and more amenable to successful models and practical computer assistance. When concerned with numerical quantities like yield curves and duration, mathematical models perform better than human judgment. On the other hand, judgment must be better since in most cases, the quantitative model outputs are guidelines to final decisions.

Just because a process is qualitative does not mean that a model and computer system cannot be built to improve that process. It is possible to build computer-based models that improve the qualitative decision-making component as well as the quantitative decision-making component.

The purpose of this chapter is to address the issues involved in building qualitative models for municipal bonds for evaluating credit quality. These models, also called knowledge-based expert systems, use concepts from logic and artificial intelligence, as well as statistics. In this chapter, we will first show why a purely quantitative approach to the credit quality problem will not be successful, and how a qualitative model addresses these shortcomings. Next, we show how a qualitative model, in the guise of an expert system, is constructed, and show what kinds of expert systems are applicable to the municipal bond credit quality problem. A particular type of system, using case-based reasoning, is then shown to be similar to the type of reasoning that analysts and portfolio managers use when making subjective judgments. Finally, we discuss an example of a system used for assessing municipal bond credit quality that utilizes all these concepts, and show the potential drawbacks of relying on a purely qualitative model.

## Quantitative Models vs. Qualitative Models

Many quantitative municipal bond portfolio models, typically based on a variation of the Capital Asset Pricing Model, have been developed in recent years. Even though the quantitative municipal bond models often out-perform analysts in some sense, the model outputs are not generally well-accepted in the day-to-day business decisions. Moreover, many analysts will simply reject the model's conclusions, especially if they are presented with binary choice between the total acceptance and total rejection of the model's decision. Some reasons for this are:

***Incompleteness of the model theory.*** The models utilize incomplete theory as well as incomplete data. For example, single index beta models are “more incomplete” than multiple index models, and both are incomplete because the betas are not absolutely known: they must be statistically estimated.

***Incompleteness of the model inputs.*** Even the best models may on occasion produce decisions much worse than a human analyst would, because some important factors may not been included in the model.

***Incompleteness of the model outputs.*** The analyst's risk preference in dealing with uncertain outcomes may differ from that of the municipal bond model. Conversely, the analyst's role is trivialized if decisions are solely provided by the model.

***Incompleteness of the explanation.*** Quantitative models provide precision at the expense of intuition and common sense.

Some analysts attempt to compensate for these limitations by making heuristic adjustments to the model in order to "tune" the results. Tuning produces a model forecast that is *consistent with intuitive expectations*, and maintains the detail and structure of the quantitative model. Needless to say, these tuning adjustments can easily be misused.

There are no formal quantitative and analytic models for municipal bond credit assessment -- that is one of the primary tasks of the analyst. Consequently, a *qualitative* or *cognitive* model of an analyst (as implemented as a knowledge-based expert system) can improve the process. For example, a formal analytical model that can be used to assess whether a municipal bond is a good credit risk may be difficult to build: the historical data will probably be incomplete for a statistical model, and the individual model factors are too dynamic. Analysts solve the problem by relying on judgment and experience. On the other hand, there are several examples of deployed expert systems that “rely on experience,” including those that perform auditing, situation assessment, compliance, and regulation.

## Types of Expert Systems

Most expert systems are developed with the following "knowledge engineering" methodology:

- Knowledge engineers investigate a task domain for system feasibility.
- Task experts are found and users are identified.
- Rules are solicited from the experts.
- A prototype is built that models the expert's task knowledge.
- An interface is built that models the user environment.
- The prototype is evaluated, modified, and deployed.

There are several types of knowledge representations that can be developed in an expert system. A rule-based system uses IF-THEN rules that can be mapped onto a decision tree: reasoning corresponds to traversing the tree. These rule-based expert system shells are based on standard *deductive* reasoning schemes, like "forward chaining" (also called "bottom-up reasoning" which is similar to the *modus ponens* of classical logic) and "backward chaining" (also called "top-down reasoning" or "reasoning by contradiction" which is similar to the *modus tollens of* classical logic).

For example, after interviewing an expert, we may derive the following rule

```
IF
        the regional economy is strong
AND     the financial conditions are positive
AND     there are no legal problems
THEN
        the credit quality is good.
```

Note that a pure deductive representation must assume either the "regional economy is strong" or not. In the real-world, this determination is subjective (and so is the conclusion that "the credit quality is good"). This means that the rule would have to be further refined if it is to be accurately used. This knowledge is "fuzzy"-- it is subjective, imprecise, incompletely specified, and seemingly inconsistent. This is one reason that in some sense, quantitative models are easier to build -- as long as the underlying statistical assumptions are justified.

However, there actually is experience in building these kinds of models. For example, most rule-based systems utilize an "uncertainty calculus" where conditions, conclusions and decision tree traversals are weighted by certainty factors. The certainty factors are then combined according to some formula which may be consistent with probability or not.

Other approaches to building a qualitative model are not dependent on logic or on a decision tree. These *inductive approaches* include:

***Case-based reasoning.*** These models represent judgments and expertise in the form of cases, not rules. Cases are represented textually and retrieved in terms of analogical reasoning strategies that are used to create database similarity indices. Given a goal case profile, the system retrieves the case “most similar” to the goal, and the “answer” for the goal profile is adapted from the answer of the most similar case in the case database. Explanations are also based on the notion of similarity and analogy. Case-based reasoning is typical of the reasoning style taught in law schools, business schools, and medical schools. In municipal bonds, this reasoning style is followed when analysts assess similarities between bonds.

***Neural Networks.*** These models also represent judgments and expertise in the form of cases, but these cases must be represented by numbers. Neural networks are similar to statistical discriminant functions (see reference 1). Neural networks are not sensitive to fuzzy data: however, like the statistical quantitative models, the number of examples that need to be provided so that the system can be "trained" is usually quite large. While convenient for pattern recognition applications, this approach has also been criticized as being too "black-box," since the internal neural network weights would have to be analyzed to provide any kind of explanation to the user.

The commonality of the inductive approaches is on the emphasis on representation of experience of training examples, cases, candidates, or episodes. In some sense, reasoning is the determination of comparing the current problem (the “case”) to past experiences. In neural network and quantitative approaches, the cases can be viewed as training patterns; in the case-based reasoning approaches, the cases may refer to actual experiences, profiles, or counter-examples.

## Case-Based Reasoning and Municipal Bond Credit Quality

Based on the availability of the data, the nature of the judgments, and the style of analyst decision-making, it would appear that case-based reasoning is the most suitable approach for a qualitative model for municipal bond credit quality. In a case-based model, the model ranks the bonds by degree of credit quality relative to the ideal credit.

Case knowledge requires the identification of the model factors for a bond profile, as well as the similarity rules that distinguish particular factor values from each other. For example, if the economic conditions of the service area is one factor, then how do we distinguish a “booming economy” from a “stagnant economy”?

Case factors can depend on other case factors as in a hierarchy. They can be qualitative as well as quantitative. One of the advantages of this approach is that qualitative and quantitative knowledge can be consistently integrated.

Case knowledge also requires rules for the combination of similarities in each factor, so that the concept of episode, example, and experience can be represented. We also need to specify the rules that define the answers: in case-based reasoning, possible answers include “the most similar cases.” If we need a single answer, we need to know how to select the best.

Similarities can be expressed as numbers between zero and one so that they can be modeled in terms of *probabilities* or *likelihoods*. Cases are evaluated to form an *analogy* with respect to the current case profile. The case with a score that has the closest *match* is the case that case-based reasoning system concludes is the best. This Selection-Analogy-Match is the central paradigm to interpreting case-based reasoning system results.

Most case-based reasoning models support different similarity functions for the model factors. The factor types are an extension to those types found in quantitative models using descriptive statistics. Factor types include

***Choice***. These values denote a *nominal value* (any distinct symbol) and can be used for qualitative and quantitative factors. Examples: *Yes, No, A, AAA*.

***Rank.*** These values denote a *Preference* defined by a list. These factors are usually used for qualitative knowledge. For example, analysts (or institutions) can be rated by the types of bonds they specialize in.

***Number.*** These values denote a *ratio of magnitudes* and can be used for qualitative and quantitative factors.

***Interval.*** These values denote a *correspondence from a number to a symbol* and are usually used for qualitative factors. Example: A coupon of 8.325% may denote "Average."

***Rating.*** These values denote a *correspondence from a symbol to a number* and can be used for qualitative and quantitative factors. This is opposite case from the interval. For example, a bond rating of AAA is "higher than" a rating of AA.

***Hierarchies.*** These values have the property of inheritance and are usually used for qualitative factors. For example, a public power bond is more similar to a transportation bond than it is to a hospital bond.

Case-based reasoning models combine different types of information in various ways to come up with a score. This combination is similar to collecting a set of *endorsements* or *credit*. Model builders can use these combination methods to mimic different reasoning styles, and to easily switch back and forth between the different styles. The information case-based reasoning system uses to construct a score for an attribute includes:

***Fuzzy Logic Combination Formulas***. Case scores can be derived by combining the case similarities in a manner consistent with probabilistic and fuzzy logic, so that scores are always between zero and one. Fuzzy logic formulas can be used to represent certain assumptions on the statistical independence or dependence of factors. For example, one expert may decide that the Economic and Financial factors are dependent; another may decide that they are independent.

***Factor Importance.*** Case scores are derived by combining the case similarities in a manner that reflects the significance or relative importance of the individual factors according to an analyst. For example, one analyst may decide that the Economic and Financial factors are equally important; another may decide that the Economic factor is much less important than the Financial factor. Analysts can also change the factor importance weights as part of a "tuning" or "what-if" analysis.

***Rapid What-If Tuning.*** In a case-based model, factor values and factor importance can be easily changed and the results instantaneously propagated: user decisions can be quickly evaluated in terms of "what-if" scenarios. Both of these activities are difficult to perform with quantitative models.

***Statistical Parameters.*** Many case-based reasoning models supports standard statistical correlation computations that can be used to assess possible relationships between different factor values and different cases similarity scores: it is possible to build a "quantitative" model on top of the case similarities. For example, we can perform a non-parametric outlier test to determine whether the first two bonds most similar to the ideal bond are in fact statistically similar or different from each other.

## A Case-Based Reasoning System for Municipal Bonds

Several case-based models have already been applied to financial decision-making applications for fixed income and equities (see references 2-7). In this section, we discuss a case-based modeling system for evaluating the credit quality of municipal bonds.

MuniCredit (created by AI Analytics, Inc. of New York) is an automated case-based system that analysts and portfolio managers use to assess the credit quality of municipal bonds based on whatever information is available to them. For example, the MuniCredit Public Power Model ranks public power bonds by degree of credit quality. The model scores each selected bond; ranks the bonds relative to the ideal credit; and shows the most important reasons for the rankings.

For example, the model provides percentage scores for the Salt River Project, Jacksonville Electric, and North Carolina EMPA; ranks them in descending order relative to the ideal credit, and shows that the most importance reason for the rankings is the spread between the relative economic strength of the Salt River and North Carolina EMPA service areas.

Analysts and portfolio managers can adapt the model to suit their individual styles of credit analysis. They can change every element of the model. The model then computes the composite effect of all of the information in the model, including statistical data and the subjective judgments of the user and any other third party experts, and shows the consequences of the analyst's judgments. The model is dynamic so that the analyst can adjust it to run sensitivity analyses or reflect rapidly-changing credit factors. The model uses a standard spreadsheet used by most analysts, so they only have to learn the unique knowledge based features in the model.

The basic elements of the public power model include the profile of the ideal credit, the number of factors, the method of comparing the bonds, and the relative degrees of importance of each factor.

***Profiling the Ideal Credit***. The analyst or portfolio manager profiles the ideal credit. For example, an analyst could describe the ideal public power credit as having a "very strong economy, very positive financial performance, very reliable and very stable operations, excellent management, and bulletproof legal protection".

***Selecting the Factors***. The analyst or portfolio manager selects the factors in the model, including number and level of detail. The number of factors can range from five to seventy-five, depending on the practicality of collecting the information.

The ability of the model to reason effectively with incomplete or partial information is particularly important in the world of municipal finance where information is far more difficult to gather than in more efficient markets like equities. One of the interesting aspects about the MuniCredit design is that it allows the analyst to reason with incomplete or partial information. In other words, MuniCredit can recommend and construct a particular case-based model with much fewer factors, based on the the risk and reward of information. This is actually another case-based model for the factors that looks at the degrees of importance for the factor, cost of getting the data, and whether the factor is qualitative or quantitative. This helps users build models cost effectively. Here, each factor is ranked by degree of cost effectiveness based upon knowledge such as the importance of the factor, the cost of collecting the information, the reliability of the information source, and a tradeoff between the specificity of qualitative and quantitative factors. Once a basic model is in place, it can be enhanced it in the ordinary course of business. The strong trend toward increased disclosure in the municipal market should make valuable new information available at reasonable cost.

Analysts can currently choose from 70 inter-related quantitative and qualitative factors organized into five primary categories and three levels of detail.

The five Level 1 factors are:

1. Economic Conditions of the Service Area
2. Financial Performance
3. Legal Protections and Risks
4. Operations Quality
5. Management Quality

The Level 2 factors for the first three Level 1 factors are:

1. Economic Conditions of the Service Area
    1.1. Employment Base Stability

1.2. Employment Base Diversity
1.3. Area Unemployment Rate
1.4. Area Income Growth
1.5. Historical Consolidated Growth Rate
1.6. Historical Population Growth
1.7. Historical Systems Connection Growth
1.8. Projected Population Growth
1.9. Projected Systems Connection Growth

2. Financial Performance
   2.1. Residential Rates per Kilowatt Hour
   2.2. Revenues Per Kilowatt Hour
   2.3. Sales Growth
   2.4. Power Growth
   2.5. Historical Debt Service Coverage
   2.6. Consolidated Balance Sheet
   2.7. Historical Fixed Charge Coverage
   2.8. Projected Fixed Charge Coverage

3. Legal Protections and Risks
   3.1. Security
   3.2. Bond Covenants
   3.3. Tax Opinion
   3.4. Contract/Litigation Rights
   3.5. Contract/Litigation Liabilities
   3.6. General Risk Management

These 22 Level 2 factors yield 75 Level 3 factors. For example, the (Level 3) factors for the (Level 2) Bond Covenants factor in the Legal Protection and Risk category are:

3. Legal Protections and Risks
   3.2. Bond Covenants
      3.2.1. Rate Covenant
      3.2.2. Additional Bonds Test
      3.3.3. Flow of Funds
      3.3.4. Debt Service Reserve Funds
      3.3.5. Renewal and Replacement Reserve Funds.

***Comparing the Bonds***. Bonds are compared in whatever is the most natural way for the analyst. Generally, qualitative factors should be used for intrinsically qualitative factors and quantitative factors for intrinsically quantitative ones. The legal factors are intrinsically qualitative. The qualitative factors enable MuniCredit to recognize the meaning of and gradations between natural language descriptions of each bond. For example, a tax opinion can range from "mainstream" to "very weak." The Financial Performance factors are intrinsically quantitative. The residential rates for each are most naturally expressed in per kilowatt-hour terms, such as $.05 per kilowatt-hour, and the projected fixed charge coverages in ratios, such as a 1.35 coverage. The Economic Conditions factors are a mixture of both. It is most natural to describe the diversity of the employment base as "diverse, very diverse, concentrated or extremely concentrated" and projected systems connection growth rates as percentages. The bottom line is that the model embeds the analyst's individual decision-making style and then shows the overall consequences of the analyst's unique methods and judgments.

***Factor Importance.*** The relative degrees of importance of the factors can be changed so that the MuniCredit model always prioritizes whatever factors the analyst considers most important at the time.

***Case-Based Models are Dynamic and Open-Ended***. The MuniCredit Public Power model is as dynamic as credit conditions. For example, due to the passage of the "National Energy Policy Act" in 1992, competition is now one of the most important credit factors for public power bonds. To immediately factor in the overall effect of the new competitive pressure, an analyst could adjust the model by, among others, (i) adding Competitive Pressure as a new factor; (ii) assigning Competitive Pressure a high degree of importance relative to most of the other factors; and (iii) adjusting upwards the degrees of importance of other factors dependent on increased competition, such as Management Quality. The quality of the issuer's management becomes more important in the new competitive environment than it was in the pre-NEPA era of monopolistic franchise areas.

## Summary

We have seen that many of the drawbacks of a quantitative model for municipal bonds can be addressed by a qualitative case-based model. However, just as a quantitative model must be tested and evaluated before use, so must the qualitative model. We conclude by discussing some of the pitfalls in prematurely accepting a qualitative model.

Qualitative models fail where there is *too much* reliance on judgment — in particular, judgments that are dynamic and difficult to verify. The fallibility of human judgment in many decision-making domains echoes the experience of many financial expert systems. Representational failures occurring in analysts (and the qualitative expert systems that model them) include:

**Anchoring**. This is the tendency not to stray from an initial judgment even when confronted with conflicting evidence. Many analysts are reluctant to revise their opinion in light of experience. In expert systems, this is seen in the difficulty to revise default assumptions in factor values and factor degrees of importance.

**Inconsistency**. If a pair of alternatives is presented to a subject many times, successive presentations being well separated by other choices, analysts do not necessarily choose the same alternative each time. In expert systems, this is seen in the representation of fuzzy and probabilistic reasoning.

**Selectivity**. This refers to using only a portion of the information available. Human analysts make poor decisions when they must take into account a number of attributes simultaneously: decision-makers may be aware of many different factors, but it is seldom more than one or two that they consider at any one time. One effect is that experts are often influenced by irrelevant information.

**Fallacy**. This refers to the improper use of probabilistic reasoning. Common errors include conservatism (the failure to revise prior probabilities sufficiently based on new information) and calibration (the discrepancy between subjective probability and objective probability).

**Representativeness**. This refers to the focusing on how closely a hypothesis matches the most recent information to the exclusion of generally available information.

**Autonomy vs. Collaboration.** Decision-makers do not wish to turn over control of a decision entirely to a single model. Just as a decision-maker is disinclined to surrender control of a decision to a mathematical model, he would not wish to surrender control to a qualitative model.

**Availability of Expertise**. In some domains, it is not possible to create a complete model for an expert system that will produce satisfactory results: there may be no experts with sufficient knowledge. The knowledge base required would be extremely large in order to anticipate all the possible conditions and actions.

**Conflicting Expertise**. Different experts reason with different styles. They may have different operational styles (which defines what information they require and what order they prefer it); and different functional styles (which can define a preferential problem-solving strategy, like top-down, bottom up, and middle out reasoning).

# References


[1] Edward I. Altman, *Distressed Securities: Analyzing and Evaluating Market Potential and Investment Risk*, Probus Publishing Co.,1991. ISBN# 1-55738-189-5.

[2] Ross M. Miller, *Computer-Aided Financial Analysis*, Addison-Wesley, Reading, 1990.

[3] V. Dhar and A. Croker, "Knowledge-Based Decision Support in Business: Issues and a Solution," *IEEE Expert*, Vol. 3, No. 1, Spring 1988, pp. 53-62.

[4] R.S. Freedman, "AI on Wall Street," *IEEE Expert*, April 1991. Reprinted in the *Encyclopedia of Computer Science and Technology*, Volume 28. ISBN# 0-8247- 2281-7.

[5] R.S. Freedman and G. J. Stuzin. A Knowledge-Based Methodology for Tuning Analytical Models. *IEEE Transactions on Systems, Man, and Cybernetics,* Vol. 21, No. 2, March, 1991.

[6] *Proceedings of the First International Conference on Artificial Intelligence Applications on Wall Street*, (R.S.Freedman, ed.), Institute of Electrical and Electronics Engineers (IEEE), 1991. ISBN# 0-8186-2240-7.

[7] *Proceedings of the Second International Conference on Artificial Intelligence Applications on Wall Street*, (R.S.Freedman, ed.), International Association of Knowledge Engineers (IAKE), 1993. ISBN# 0-9-38801-07-4.