To Appear *in The Encyclopedia of Computer Technology, Volume 28, Marcel Dekker, Inc., New York 1993,* pp. 15-25.

AI on Wall Street

Roy S. Freedman
Polytechnic University
Inductive Solutions, Inc.

Abstract

"Wall Street" refers to the complex, data intensive and international environment that supports the buying and selling of financial products and commodities. For the past several years, many Wall Street firms have developed systems to demonstrate how AI technology can support this environment. Many of these highly proprietary systems have been quietly deployed; others have been abandoned as incomplete prototypes. The goal of this article is to show why some Wall Street AI applications are successful and why others fail.

Introduction

He cursed Theocritus and Homer,
in Adam Smith was his diploma;
our deep economist had got
the gift of recognizing what
a nation's wealth is, what augments it,
and how a country lives, and why
it needs no gold if a supply
of *simple product* supplements it.
His father failed to understand
and took a mortgage on his land.

Alexander Pushkin, *Eugene Onegin* 1,VII (1829)
(translated by Charles Johnston)

"Wall Street" refers to the complex, data intensive and international environment that supports the buying and selling of financial products and commodities.

This environment evolved with the explosion of international trade that followed the discovery of the new world. Centralized financial "exchanges" appeared in Antwerp (1530), Hamburg (1558), London (the Royal Exchange in 1565), Amsterdam (1613), Japan (with the Mitsui family in 1673), and Vienna (1753). These financial exchanges helped organizations raise money for trade and investment, and redefined concepts such as ownership, risk, and trust in terms of financial products such as securities (stocks, bonds, and mortgages) and other instruments (such as currencies). In fact, the term "share" of stock seems to have been first used by the Dutch East India Company in 1609 during the raising of capital for their New Amsterdam enterprises [1].

New York City had an auction market that supported the trade of financial products as early as 1725. This market initially supported the trading of commodities (wheat and tobacco) and later expanded to support the trading of Revolutionary War Bonds and other securities. In 1792, several brokers separated from the auction market to form a new market which specialized in securities. They initially met under a buttonwood tree at 68 Wall Street. This organization later became the New York Stock Exchange, an organization virtually synonymous with "Wall Street."

Market participants know that improved information technology provides competitive business advantages, and improves the fairness of the market for all participants. A market is fair when all market participants have access to the same information about market products. In the nineteenth century, market information was publicly displayed on a large blackboard (this was called the "Big Board"at the New York Stock Exchange). Like blackboard expert system models, market participants would all look at the board and fire rules to determine their market activity. In the Wall Street environment, this blackboard model proved inadequate: as the number of products increased, the board quickly ran out of space. Moreover, information was not really public: it was public to those market participants who had access to the board.

Public information access (and increased market fairness) was improved with the invention of the *ticker*-- an electronic device that can display information on recent trades to *any* participant. The stock ticker was invented by E. Calahan in 1867 (an employee at the New York Stock Exchange), and was later enhanced and modified by Thomas Edison.

Today, the relationship between improved information technology and market fairness and competition is still valid. The ticker still exists, and the principal of disseminating and exploiting market knowledge has been seen in the deployment of several systems utilizing expert system and artificial intelligence techniques. Many Wall Street firms have developed systems to demonstrate how this technology supports the buying and selling of financial products and commodities in applications such as automated news understanding, credit assessment, risk management, market surveillance, auditing, and trading [2,3]. Many of these highly proprietary systems have been quietly deployed; others have been abandoned as incomplete prototypes.

The goal of this article is to show why some Wall Street AI applications are successful and why others fail. To do this, we first model the Wall Street problem domain from a knowledge-based perspective. We provide two declarative representations of the Wall Street environment: an object oriented view and a process view of the trading and compliance domain. Both models demonstrate the complexity inherent in formally defining the Wall Street environment. This complexity prevents a complete predictive analytical model of the Wall Street environment, and leads us to discuss some of the reasons for an expert systems approach. This approach is motivated by situations where a cognitive model of human judgment performs better than an analytical model. Since human judgments are dynamic and are often difficult to test in predictive situations, we expect that the expert system approach will fail if the underlying cognitive model is inadequate. Finally, we discuss a knowledge-based approach that does seem to integrate the best of both approaches: the development of collaborative systems.

Wall Street: An Object-Oriented View

Object-oriented models seek to simultaneously specify objects together with morphisms (rules and operations) defined on these objects. Objects send messages to other objects; object semantics are defined by the reception of messages.

Our object-oriented view of the Wall Street environment defines a complex set of competitive worlds that are highly distributed in space and time. Each world is an instantiation of markets, market participants, products, orders, news, and regulatory structures. Figure 1 summarizes some of the objects and messages in a world in the form of a semantic network.

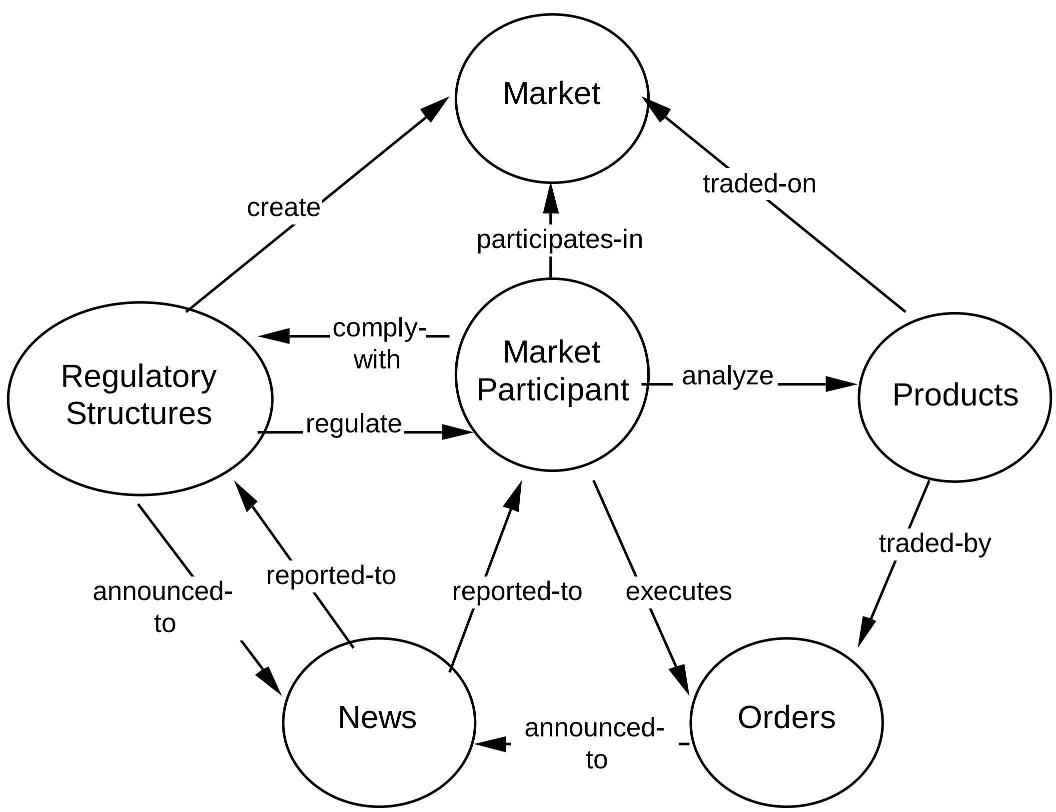

Figure 1. A Semantic Network Representation of a Wall Street World

Products represent assets (e.g., securities such as bonds and stocks) or contracts (e.g., options to buy or sell an asset within a specified future time), or may represent an index computed on an underlying product (e.g., an interest rate, or an average price of a "basket" or set of securities).

Market participants buy and sell products by performing a set of actions that maximize an economic objective function. Buy and sell transactions are recorded as orders.

One way that market participants can maximize an economic objective function is by firing the following rules:

> Buy low and sell high.
> Do it before anyone else.

If these actions are common to a set of market participants, then the only market participants that will maximize their economic objective functions are those that have a *product knowledge advantage* (so that they can effectively determine what is "high" and "low"), and those that have an *order execution advantage* (so that they do not wait to achieve these actions: waiting may change the price to a sub-optimal level).

Orders are used to indicate different timing and size parameters for trades. There are many different types of orders that are based on the number of combinations for size (product quantity), side (buy or sell, long or short), price limits, time limits, and other constraints (stop orders, orders on the opening or close, "fire or kill" orders, and others).

The improvement in order execution has been continuing since the installation of telegraph communications over one hundred years ago (and continues today in the context of global markets and 24-hour trading). These improvements include better order, billing and clearance systems that can handle large numbers of transactions. In this context, AI has been applied to financial telecommunication network management; funds transfer; and image, voice, and forms processing.

Product knowledge advantages have been traditionally acquired through the efforts in fundamental and technical financial analysis. Fundamental analysis refers to evaluating a product with respect to deep knowledge about assets and risks and classical financial ratios. Technical analysis refers to evaluating a product based on the shallow knowledge of price and volume trends. Trends may be statistical (based on so-called "moving averages") or non-statistical (based on ad hoc theories).

One goal in the Wall Street environment is to discover product knowledge advantages in order to assess the price of a product. If all market participants have perfect information and the same execution capabilities, then *on the average*, none will have any knowledge advantage over the other. In this ideal environment, the price of a product on a market converges to its "true price."

In the real Wall Street environment, there are price deviations between the true price and the market price because not all traders have complete information. These deviations are further exaggerated by rumors, "good news," "bad news," panics, and crashes. The existence of the true price is an consequence of the so-called efficient market hypothesis: the true price of an object is the price observed if all traders act rationally, have perfect information about the object, and are all participating in the market. This hypothesis is usually accepted when all traders have access to identical telecommunication and information systems: all information obtainable about an object is known to all participants. Consequently, in this rational and perfect information scenario, there is no profitable way to speculate.

Some knowledge advantages may result from acting on knowledge acquired from non-public information sources. Such *insider trading*-- the use of proprietary or non-public “inside” information for financial gain-- is not legal in many markets.

Regulatory structures prevent illegal knowledge advantages in the market so that trading remains fair and orderly. Regulations are laws that exist in at the government agency level, the market level, and the market participant compliance level. In some sense, the regulatory structures define the market. In our model, regulatory structures monitor trading and may execute particular actions if they determine that the market is not fair or orderly. Regulatory actions include the temporary halting of trading of an object in a market, or the permanent prevention of trading of a market participant, or the temporary halting of trading of all products in a market. Other actions include the levying financial fines and penalties for rule violations.

One of the more challenging applications of AI on Wall Street is in the detection of (legal and illegal) knowledge advantages in trading and regulation. Such systems have been deployed for assessing potential opportunities (risk assessment, underwriting, exposure), evaluating lost opportunities (auditing), and compliance with different regulatory structures.

Markets comply with different regulatory structures. The most regulated markets are *exchanges*. These are centralized markets registered with a government agency, such as the Securities and Exchange Commission (SEC). The least regulated are *crossing markets*, where buyers and sellers agree between themselves on the order. Orders that do not take place on an exchange but do take place within a system of brokers are said to take place on over-the-counter (OTC) markets.

Different markets support different types of objects. Some are market makers (or dealers); their regulated function is to assist in the trading of products quickly (so that the market is *continuous*) and to insure that the deviation between the true price and the market price of an object is small (so that the market is *liquid*). Dealers provide market continuity and liquidity by keeping an inventory of the traded object.

Others market participants are brokers: their regulated function is to represent traders in an order (for a commission) and match buyers and sellers. Market makers who are simultaneously brokers and dealers are called specialists. Because of this dual role, specialists must comply with very strong market rules. For example, a specialist is prohibited from selling a product if, on the the previous executed order, the price of the product went down. This action of “selling on a down tick” is prohibited because it may affect the market by lowering the price even further--in a potentially destabilizing way-- thus affecting market liquidity.

To complete the above object-oriented specification of the Wall Street environment, we require the existence of several inter-dependent markets. This linkage can be accomplished in several ways

1. A product can be traded on several markets at once.
2. A product trading on one market can effect the trading of a different product on another market.
3. A market participant can participate on several markets at once.
4. Information may (or may not) reach all market participants.

The first type of linkage has been in existence for at least one hundred years. This corresponds to the case where shares of ACME Corporation trades simultaneously on the New York and London Stock Exchanges.

The second type of linkage may be explicit or implicit. For example, market participants may trade contracts that enable them to have the right to buy a product (the "underlying" product) at a particular price during a stated time period on an an *option* market. With suitable statistical assumptions, it can be shown that the price of the option is explicitly mathematically related to the market price of the underlying object. A *future* is another example of a "derivative" object: a future is a commitment to buy or sell an underlying product at a particular price during a stated time period. There can also be options on futures, futures on indices ("indexed" futures), and so on.

The third type of linkage is seen in the practice of trading on multiple markets, either to take advantage of some "arbitrage" opportunities, or to manage risks by hedging. In general, their may only be implicit relationships between the prices different products on different markets; precise mathematical models detailing explicit relationships may not be adequate for predicting prices. For example, a trader may believe that the price of yen is dependent on the price of gold, but may base this belief on judgment and experience, and not on any formal analytical econometric model. The interplay between analytical models of Wall Street, human judgment, and AI techniques will be discussed later.

The fourth type of linkage introduces the complexity of information flows. In general, not all market participants are totally connected to the information available in each market.

In our model, the object used to disseminate information among market participants and regulatory structures is called ***News***.

Figure 1 is a model of a single world: a more complete model of Wall Street would consist of instances of the figure (which can be drawn in multiple layers), with inter- and intra- layer messages linking the different markets, products, news, market participants, and regulatory structures in different layers.

Any causal model of the Wall Street environment must be incomplete. As market participants become aware of the different information flows affecting the environment, feedback from news causes the objects to change and evolve new products, new markets, new regulatory structures, and new market participants.
Any complete causal analytic models either do not exist or are too complex to justify short-term predictions. In this sense, the Wall Street environment behaves like a living organism.

Wall Street: A Process View

Process models emphasize information transformation. In this section, we model the Wall Street environment in terms of a generic structure of processes that create and monitor events. Our model is a variation of the Predict-Detect-Isolate-Correct feedback loop discussed in the ISO network management model. Network management is another domain known for event complexity and AI applicability [4].

A procedural model can be used to model a trading and regulation situation. Both are shown pictorially in Figure 2. The difference between these scenarios concern their dual interpretation of an event. In trading, the goal is to discover events (with low probability) that represent a buy or sell opportunity; in regulation, the goal is to discover events (with low probability) that represent a possible indication of something unusual that must be further investigated for non-compliance.

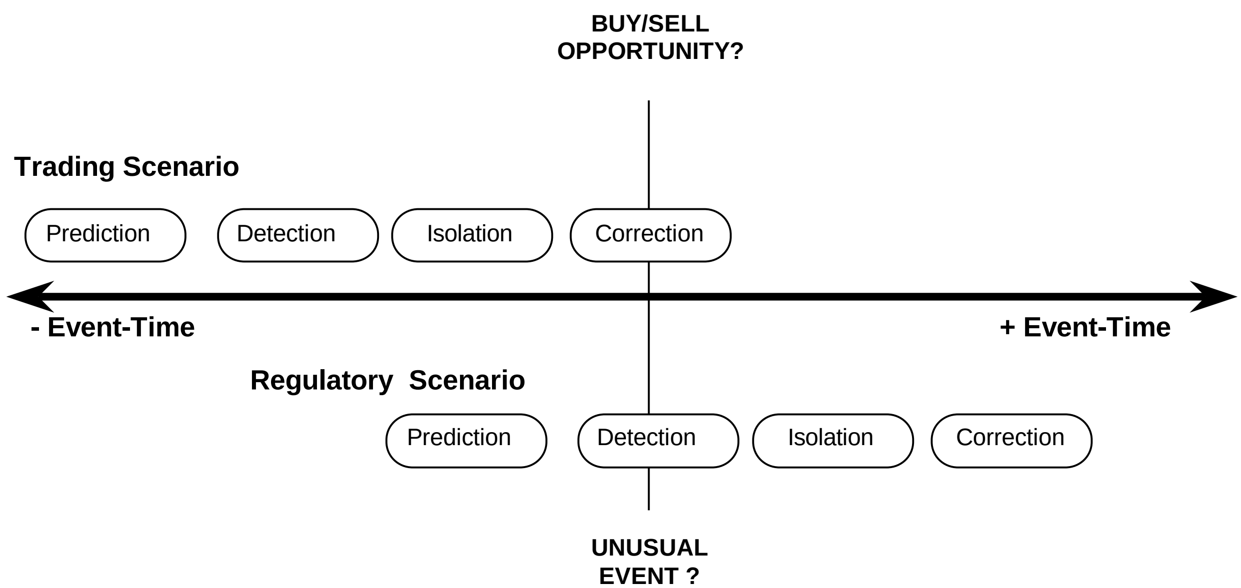

Figure 2. Trading and Regulation Scenarios.

The basic Wall Street model is shown as a data-flow diagram in Figure 3. In this diagram, we assume the efficient market hypothesis that that all world knowledge (market and news) is available to each process.

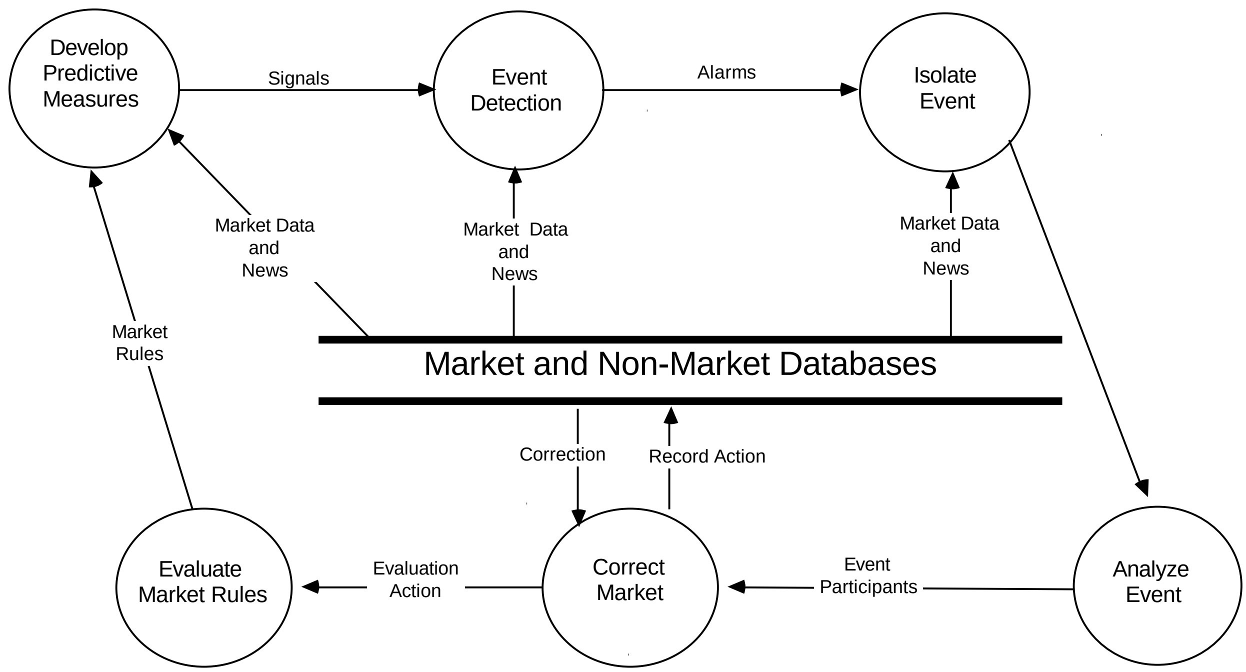

Figure 3. Wall Street Environment as a Real-Time Process Model

Prediction refers to the determination that an event will occur in the future. Prediction involves the creation of a model that can represent past experiences.
In a market trading scenario, the prediction process determines the likely buy or sell signals by defining a computation based on certain database inputs. For example, a trading model may have a set of predictive rules that are based on price and volume movements, such as "If the 10-day moving price average crosses the price curve in a broadening bottom pattern, then signal a buy, otherwise hold."

The behavior of the market regulation scenario is similar. Here, the prediction process determines the likely signals that indicate that an unusual event has occurred. For example, a model may have a set of predictive rules that are also based on price and volume movements, such as "If the price of a product increases by 15%, then it is likely that there was a unusual event about the product."

Detection refers to the determination that an event has already occurred. Detection is concerned with discovering patterns. In a market trading scenario, the detection process determines that a buy or sell opportunity has occurred by firing the rules on the available data. In a market regulation scenario, the detection process determines that an unusual event has occurred. Rule firings trigger alarms (or alerts).

Isolation refers to the process of analyzing the event to determine event components or event participants. Given a pattern, isolation is concerned with discovering component patterns. For example, in a market trading scenario, the isolation process can analyze several rule firings (alarms) and extract the most likely set of products within on which to Buy or Sell. These objects may also involve derivative objects.

Correction refers to the enforcement of rules adjust the predictive model. Correction is concerned with executing an action (such as producing a plan or explanation) that may change the world or change the model. In a market trading scenario, one output of this process can be an order for an actual trade. In a market regulation scenario, outputs can be trading halts for products, for participants, or for the entire market. The correction process may maintain a database of past actions and predictions that can be used to update the predictive model.

There is a relationship between the process model and the object-oriented model. In some sense, all objects interact with their world by performing the Predict-Detect-Isolate-Correct loop. The process model helps form a local view of the world. Because of this local view, we can scope small domains to assess the applicability of different AI technologies.

Some Wall Street applications (such as credit screening) represent knowledge by some standard stable decision procedures. In these applications, deductive methods such as those based on production rules provide the easiest way to build predictive cognitive models. Probabilistic reasoning can be embedded in these models fairly easily. The predictive model is based on a search that is constrained by the rules. Here, the Predict-Detect-Isolate-Correct loop can correspond to the steps (i) build the rule base; (ii) fire rules on global data structure; (iii) fire rules on detailed data structure; (iv) perform the recommended system action (and possibly create new rules). This last step is extremely important to many Wall Street applications because it addresses the non-monotonic behavior of the world: information acquired at a future time may be inconsistent with default assumptions. In deductive systems, such consistency is required.

Other Wall Street applications cannot represent knowledge by a standard decision procedure. In these applications, it may be appropriate to build predictive cognitive models using non-deductive (or inductive) methods. These methods do not require logical consistency.

Case-based approaches are appropriate when the application requires the representation of actual experiences. In many applications, this necessitates the creation and maintenance of a database of exemplars. The predictive model is based on searching for experiences that are "most similar" to the current situation, and adapting those experiences to fit the current situation. Here, the Predict-Detect-Isolate-Correct loop can correspond to the steps (i) build and index the case base; (ii) find the most similar exemplars; (iii) select the best exemplar and adapt it; (iv) perform the recommended system action (and possibly create new cases).

Machine learning techniques may be appropriate when a database of actual experiences exists and a standard decision procedure *is assumed to exist*. Here, part of the prediction process consists of inferring rules that can be used in a deductive Predict-Detect-Isolate-Correct loop.

Connectionist (neural network) approaches are appropriate when the application requires the representation of actual experiences in a "compiled form." This approach is subsymbolic: it rules on a numerical representation and computational models utilizing linear and non-linear discriminants. In some sense, this approach can also be considered similar to the analytic approach: the output of the prediction process is can be considered to be an analytic (statistical) model.

From an operational perspective, the two most challenging problems concerning the building of systems based on the above process models for Wall Street applications are complexity ("can it work?") and validation ("does it work?"). Successfully deployed Wall Street systems address both problems.

The complexity of Wall Street systems mimic the complexity of the environment. This complexity can be reduced to understandable levels by utilizing system components that already exist. To do this, we require standard definitions for user and data interfaces to permit the interoperability between implementations of cognitive models (knowledge-based systems), implementations of analytic models, and users.

Validation is as much a function of management and user commitment to the system as system complexity: even if complexity is well managed, a system that cannot be validated by demonstrating the support of specific business and user goals will not be deployed.

Expert Systems on Wall Street: Analytical Models vs. Cognitive Models

Many analytical models have been developed for the Wall Street environment in recent years. Several of these models have been used successfully for decision-making and prediction, such as the well-known Capital Asset Pricing Model and the Black-Scholes model for options pricing [5].

In other domains (such as economic forecasting), even though some analytic models often out-perform human decision-makers, the model outputs are not generally well-accepted. Many analysts will usually 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.

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 have 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 model. Conversely, the analyst's role is trivialized if decisions are solely provided by the model.

Incompleteness of the explanation. 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 analytic model. As Pindyck and Rubinfeld [6] observe:

> There is a ...method that is often used to make minor adjustments in ... models, particularly those that are used for forecasting purposes. This method is called 'tuning' and consists of making small changes in some of the model's coefficients,...so as to improve the ability of the model to forecast.
> [Tuning has] come to be used in large... forecasting models, particularly those constructed for commercial or business applications (often they are adjusted to keep the forecast "in line" with intuitive forecasts - thus to some extent negating the predictions of the model). Needless to say, [these adjustments] can easily be misused (and often are).

The Wall Street environment is rich enough to permit domains where human judgment performs better than analytical models. Consequently, a cognitive model of an analyst (as implemented as an expert system) may perform better at predictive tasks then an analytic model. For example, a formal analytical model that can be used to assess whether an individual is a good risk on a loan is difficult to build. This problem domain relies on judgment and experience. On the other hand, there are several examples of expert systems for this type of screening application. Other Wall Street domains that have incomplete analytical models but rely on experience include auditing, clearing, compliance, and applications involving natural language processing (such as the parsing of news and market data).

Cognitive models fail for domains where there is *too much* reliance on judgment: in these domains, judgments are dynamic and their representations are difficult to quantify and verify.

The fallibility of human judgment in many decision-making domains echos the experience of many financial expert systems. Representational failures occurring in human analysts (and the expert systems that model them) include:

Anchoring. This is the tendency not to stray from an initial judgment even when confronted with conflicting evidence. Humans are reluctant to revise their opinion in light of experience. In expert systems, this is seen in the difficulty to revise default assumptions in non-monotonic reasoning.

Inconsistency. Given quantities A, B, and C, consistent behavior would require a subject to treat them as though they satisfied the following two properties:

1. Exclusivity of comparison.
   Either A > B *or* A < B *or* A = B.

2. Transitivity of comparison.
   If A > B *and* B > C *then* A > C.

Humans tend to violate both properties: if a pair of alternatives is presented to a subject many times, successive presentations being well separated by other choices, a given subject does not necessarily choose the same alternative each time. In expert systems, this is seen in the representation of fuzzy and probabilistic reasoning. In these systems, if proposition A has certainty p_A of occurring, then in some systems, it is consistent to infer that the certainty of the proposition "A and not-A" has certainty $p_A*(1- p_A)$.

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 [6]. 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 cognitive 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).

The utilization of analytical models, human judgment, and expert systems on Wall Street is not necessarily disjoint. Inadequacy of the different models can be improved by not only by integration, but by building *collaborating* models.

Collaborative Systems

The desirability of developing techniques by which humans and computers collaborate in making decisions, rather than the decision being made by one or the other, has been recognized for some time. As early as 1961, Yntema and Torgerson [7] questioned how to combine the speed of the computer with the "good sense" of the human user, without sacrificing too much of either. They proposed to let the machine make the decisions according to simple rules, but require the analyst to monitor the result and change the machine's answer if the analyst finds the results too foolish.

Analytic models have become much more complex since 1961. However, all abstract models are only approximations: optimization achieved with respect to the model is not the same as optimization with respect to the real Wall Street environment. Ultimately, only a human can judge if the discrepancy between the real Wall Street and the model is large or small.

The combination of analytical models, human judgment, and cognitive models leads to a new kind of intelligent tool. They are not so much intelligent assistants as they are *collaborators.*

From an operational perspective, a collaborator should relieve an analyst of routine computation, data handling, and decision-making. When requested, a collaborator should also explain its reasoning. In this regard, collaborators are similar to apprentices and a tutors in that they compare user behavior to expert behavior and attempt to minimize the difference by negotiation. The objective of the negotiating process is to influence the user's behavior, making it as rational as possible from the perspective of the domain expert knowledge.

A general program for an "artificial laboratory" of such tools was also proposed in [8], where the assistance provided analysts was classified into three components: *model developers and representers*; *model testers*; and *model refiners*. Such tools may be used as collaborators in discovery (model creation) as well as collaborators in model utilization.

Relationships between analytical models, human judgment, and cognitive models are shown in Figure 4. Some examples of Knowledge-based collaboration between expert systems and analytic models for the Wall Street environment are discussed in [10] and [11]. Other examples of intelligent computational tools that assist in the development and representation of analytic models are discussed in [9]. Knowledge-based tuning (in the context of of econometric models) is discussed in [12]. Some examples of knowledge-based expert systems (based on cognitive models of financial experts) are discussed in [2,3].

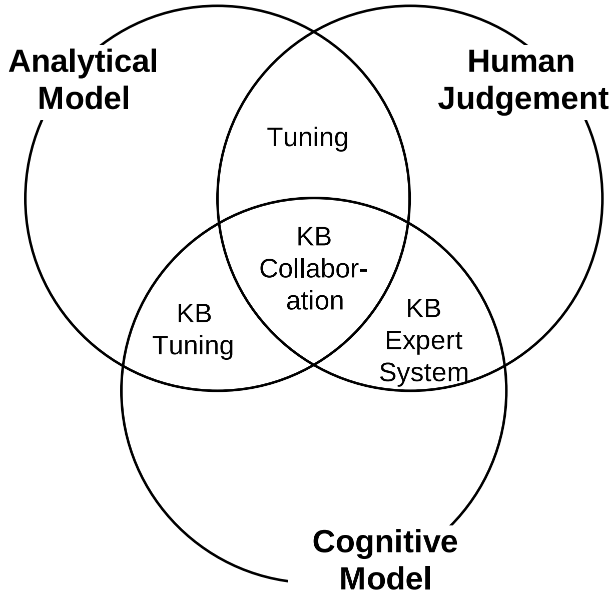

Figure 4. Analytic Models vs. Cognitive Models vs. Human Judgment

Summary

We have reviewed the Wall Street problem domain from a knowledge-based perspective by providing a global object-oriented view and a local process view of the trading and compliance domain. Both views demonstrate the complexity inherent in formally defining the Wall Street environment. This complexity prevents a complete predictive analytical and cognitive model of the Wall Street environment.

AI technology is most applicable to problems where a cognitive model of human judgment performs better than a mathematical model. Essentially, we have shown that the reason why some Wall Street AI applications are successful and why others fail is due to the representational adequacy of the underlying cognitive model. We also discussed how integration and validation have an impact on system deployment.

Wall Street continues to be a challenging domain for innovative AI applications. Today, successfully deployed AI systems have the important operational property of being successfully integrated with other systems. Future Wall Street AI systems will continue to build on the advantages of system integration by demonstrating the intelligent collaboration among different components.

References

1. Bernard Grun, *Timetables of History*, Simon & Schuster, New York, 1979.

2. David Leinweber, "Knowledge-Based Systems for Financial Applications," *IEEE Expert*, Vol. 3, No. 3, Fall 1988, pp. 18-31.

3. Ivy Schmerken, "Wall Street's Elusive Goal: Computers that Think Like Pros," *Wall Street Computer Review*, Vol. 7, No. 9, June 1990.

4. C. Joseph, A. Sherzer, K. Muralidhar. Knowledge Based Fault Management for OSI Networks. *Proceedings of the Third International Conference on Industrial and Engineering Applications of Artificial Intelligence*, ACM, New York, 1990, pp. 61-69.

5. Ross M. Miller, *Computer-Aided Financial Analysis,* Addison-Wesley, Reading, 1990.

6. R. S. Pindyck and D. L. Rubinfeld, *Econometric Models and Economic Forecasts*, McGraw-Hill, 1976.

7. D. B. Yntema and W. S. Torgerson, "Man-Computer Cooperation in Decisions Requiring Common Sense", *IRE Transactions on Human Factors in Electronics*, HFE-2 (1961), pp 20 - 26. Reprinted in W. Edwards and A. Tversky, *Decision Making*, Penguin Books, 1967.

8. M. Lacy, "Artificial Laboratories," *AI Magazine*, Vol. 10, No. 2, 1989, pp. 43-48.

9. H. Abelson, M. Eisenberg, M. Halfant, J. Katznelson, E. Sacks, G. Sussman, J. Wisdom, K. Yip, "Intelligence in Scientific Computing," *Communications of the ACM*, Vol. 32, No. 5, May 1989, pp. 546-561.

10. C. Apte and R. Dionne, "Building numerical sensitivity analysis using a knowledge based approach," *IEEE Expert*, Vol. 3, No. 2, Summer 1988, pp. 371-378.

11. 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.

12. 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 (forthcoming).