

**AN INTELLIGENT ASSISTANT
FOR FINANCIAL HEDGING**

by

Michel Benaroch

Leonard N. Stern School of Business
New York University
New York, NY 10003

and

Vasant Dhar

Leonard N. Stern School of Business
New York University
New York, NY 10003

October 1991

Center for Research on Information Systems
Information Systems Department
Leonard N. Stern School of Business
New York University

Working Paper Series

STERN IS-91-33

This paper appeared in the 5th International Conference on Artificial Intelligence Applications, Miami, March 1990.

An Intelligent Assistant for Financial Hedging

Michel Benaroch

Vasant Dhar

Information Systems Department
New York University
New York, NY 10006

Abstract

Problems in Finance, particularly those involving risk assessment and management, have been slow to yield to expert systems technology for two reasons. First, expert reasoning in such problems is often based on "first principles" instead of "situation-action" rules that characterize most expert systems. Secondly, the knowledge involved, such as that about financial instruments, is constantly changing. This would make it extremely difficult to keep a rule-base accurate. We have developed a representation in the domain of financial hedging that has the following characteristics. First, it allows for reasoning qualitatively based on first principles using the fundamental quantitative valuation models that characterize each instrument. Secondly, it uses object oriented concepts and inheritance to minimize the effort needed to set up the knowledge base and keep it current. Thirdly, it includes a calculus for derivation of qualitative knowledge of "one-dimensional-order", which allows it to solve problems where optimality constraints are qualitative.

AI Topic: Knowledge Representation, Qualitative Reasoning.

Domain Area: Financial Hedging, Risk Management.

Language/Toll: C++, Common Lisp, IBM-AT.

Status: Prototype fully implemented and currently under testing.

Effort: 1.2 person years.

1 Introduction

Today's increasing globalization of markets and sophistication of financial instruments is increasing the complexity and the number of ways to control risk. To detect opportunities, however, what is required is an intelligent interpretation of the metrics that characterize each instrument.

Our interest in risk management is on the design of

hedge vehicles that allow us to control the balance between risk and reward. This balance is a function of the instruments available in the market, and the trader's preferences and beliefs. Accordingly, the problem can be characterized by the following features:

(1) the combinations associated with considering all hedging alternatives is explosive. Traders, however, consider only a subset of alternatives based on their specialization in a few types of instruments.

(2) instruments are also evaluated along qualitative-valued attributes (e.g., *liquidity*) describing desired characteristics of hedge vehicles. Conceptually, the problem of hedge design can be viewed as a multi-objective optimization problem, which can be shown to be NP-complete (and hard to solve to optimality). Good heuristics are therefore important.

(3) the experienced trader invariably wants the hedge vehicle to take into account his assessments of how certain economic factors are likely to behave over a specific time period. Many of these assessments are expressed qualitatively based on current market information or historical trends. In order to derive a payoff from such assessments, the trader must underwrite his own "internal" option strategy, one that accounts for what he is willing to pay and what he is willing to risk under various market scenarios.

These characteristics encourage the use of expert systems technology in the design of an intelligent assistant for hedging. In particular, the role of such an assistant is one of designing hedge vehicles based on the trader's specification of an event to be hedged and of his preferences and beliefs. It is also desirable that the system be able to explain its recommendations.

Secondly, since new instruments are periodically introduced into the marketplace and old ones are discarded, it is important that the system require minimal effort for knowledge base updating. A representation that minimizes such effort is therefore important.

This paper describes the knowledge representation we have used to develop a prototype expert system for finan-

cial hedging. This representation captures the “deep domain knowledge” that experts use to reason about hedging decisions. It is flexible enough to reason in terms of the basic principles of risk assessment – something good traders are required to perform in order to make good decisions.

The paper is organized as follows. Section 2 reviews the domain of hedging. Section 3 explains the need for reasoning from first principles in hedging and describes the representation we use to facilitate such reasoning. Section 4 explains how this representation is used to design hedge vehicles. Finally, Section 5 presents the current limitations of our prototype.

2 The Domain

The objective of hedging is to offset potential losses due to possible future uncertain events. The underlying principle of hedging is, given an uncertain event and an asset that is sensitive to that event, to match that asset with a liability (asset) whose sensitivity to the event is similar (opposing).

Consider for example a firm that plans to issue bonds in order to raise capital. If interest rates rise before the date of issuance, the firm will have to offer a higher rate on the bond, which raise the cost of issuance. Suppose the firm believes that interest rates are likely to increase prior to the issuance date, but there is a small possibility that they will decline. The firm could therefore define a “cap issuance rate” protection profile (i.e., the firm wants to protect against an increase, while preserving the ability to benefit from a decline, in interest rates).

One hedge vehicle the firm could use to cap its issuance rate is to purchase put options on bonds with similar maturity and price sensitivity to that of the debt issuance. A put buyer has the right, but not the obligation, to sell to the put seller a bond at a specified exercise price at some future expiration date. A rise in interest rates causes bond prices to decline, giving the put buyer an opportunity to gain by exercising his option to sell bonds at the higher agreed price, such that the gain on the sale offsets the extra cost of issuing bonds at a higher yield rate. Alternately, a decline in interest rates will cause the bond price to decline, rendering the put option worthless; however the gain from issuing bonds at a lower rate will offset, and more, the cost of buying the put option.

Hedge design involves thousands of instruments, and requires consideration of new instruments that are constantly introduced to the market (e.g., newly issued bonds) and elimination of instruments that can no longer be used (e.g., matured bonds). Hedge instruments can be generic (e.g., an exchange traded option), compounded (e.g., strangle – sell a Put option and a Call

option), or synthetic (e.g., an option replica constructed from cash and that option’s underlying security).

In practice, traders will always prefer one hedge vehicle over others based on both quantitative-valued attributes such as *vehicle-cost* and qualitative-valued attributes such as *setup-complexity*. Overall, the design of hedge vehicles is influenced by the event to be hedged, the current and the projected state of capital markets, the hedger objectives and constraints, and the instruments available on the market.

3 Knowledge Representation

While rules may be useful as summary retrospectives of expert behavior, they are often not flexible enough to generate even simple reasoning in novel situations [2]. The following factors suggest why a rule-based representation is unlikely to work, and indicate the type of representation that is more appropriate:

(1) Combinatorial explosion in search space: enumerating all possible combinations of values of economic factors and associating them with actions would require experts to envision all possible situation-action rules, something that makes the problem hard to solve, and also infeasible from a knowledge acquisition standpoint. What is even more difficult is refining the rules as inadequacies in the knowledge-base are recognized.

(2) Limited applicability of knowledge: rules are usually compiled in the context of one specific problem area. Since such knowledge is derived using knowledge about basic principles that apply across similar problem areas (basic relationships between domain entities that are parameterized differently according to the specific problem and situation), it should be usable in reasoning about related problems. This essentially requires a representation that refers directly to the basic structural entities of the domain.

(3) Closed knowledge base: it is difficult in a rule-based system to integrate new knowledge, such as about new instruments. To be truly open, a knowledge base must be able to relate easily new domain entities to the ones it already knows about.

We feel that an intelligent assistant for hedging must use a representation that is flexible enough to reason in terms of the basic principles of risk assessment (basic economic factors and relationships). It must capture the deep domain knowledge that experts use to reason about hedging decisions involving both quantitative and qualitative factors. In fact, it should allow for the derivation of rules one might elicit from an expert if one were trying to develop a rule-based system.

The rest of this section describes the representation underlying our system. This representation consists of

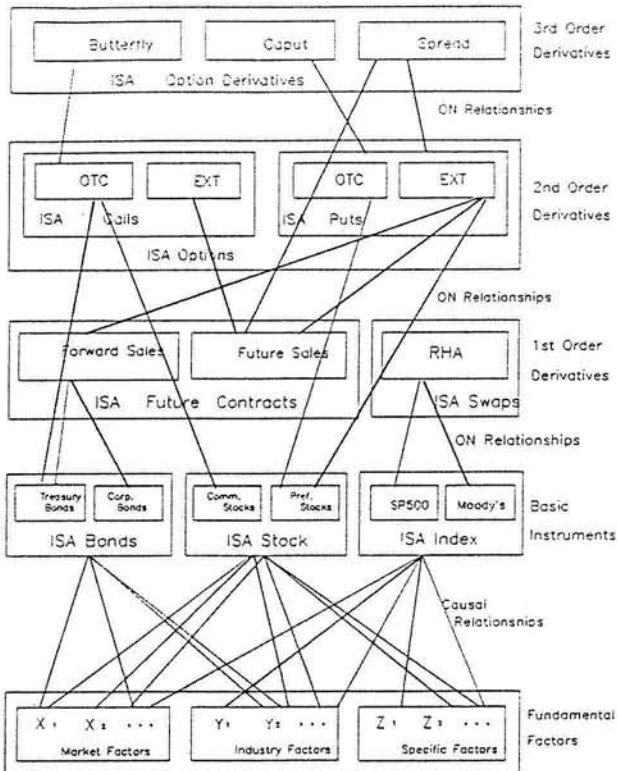


Figure 1: Relationships between fundamental factors, basic instruments, and derivative instruments

four architectural components: (1) an ISA hierarchy that reflects specialization relationships between instruments; (2) a hierarchy that reflects classes of causal relationships between instruments and fundamental economic factors; (3) a causal reasoning engine for modeling the behavior of instruments' attributes under different behaviors of economic factors; and (4) a calculus for combining what is often referred to as "qualitative knowledge of one-dimensional order".

3.1 The ON and ISA Hierarchies

Much of the knowledge involved in hedge design centers around generic financial instruments and fundamental economic factors that determine their behavior (i.e., monetary-value, annual-return, risk). Generic instruments can be classified as basic instruments (e.g., bonds) or derivative instruments (e.g., options ON bonds). The behavior of basic instruments is determined by the behavior of fundamental factors, while the behavior of derivative instruments is, in addition, determined by the behavior of lower level derivatives and of basic underlying instruments (see Figure 1). We call the hierarchy that reflects this knowledge the *ON hierarchy*.

A comparison of attributes of instruments that experts

use to reason about hedging decisions shows many similarities among instruments. In fact, instruments are also related via specialization relationships. For example, future sales and forward sales are both types of future contracts. Figure 1 shows the *ISA hierarchy* reflecting this knowledge.

The sensitivity of an instrument to a specific fundamental factor is considered as a causal relationship (e.g., an increase in interest rates causes bond prices to decline). It is modeled using a structural equation. The analytical solution of a set of structural equations describing the major causal relationships between an instrument and fundamental factors is called the *valuation model* of that instrument [3]. It is the mathematical model used to price that instrument. Instruments of the same class have the same valuation model.

We have implemented these two hierarchies as an object-oriented network to which we shall refer to as the *ON-ISA network*. Each node in the network represents a generic instrument or a class of generic instruments. It stores knowledge about attributes of the individual or class represented (e.g., *maturity-date* and *trading-volume*). The knowledge stored in each node is determined by the kind of reasoning required with that node in the hedge design process. The ON and ISA relationships are represented using physical links between nodes. The relationships between every class of instruments and fundamental factors are represented using the set of structural equations whose solution is the valuation model of that class.

3.2 The Qualitative Causal Reasoning Engine

An analysis of guides to hedging (e.g., [1], [4], [7]) reveals one type of "compiled rule" experts use to identify instruments that provide the trader's protection profile (payoff profile the hedger wants to derive). For example, "If the firm's objective is to cap debt issuance rates and the firm's issuance plan is highly uncertain, then consider buying a Caput (Call on Put)."

Experts justify such rules by doing a qualitative causal analysis using knowledge about structural relationships between fundamental factors and instruments. For example, "If interest rates rise, the Put will increase in value, causing the Caput to become more valuable. Thus, the gain on the Caput offsets the extra cost of issuing at higher rates. If interest rates decline slightly, the Put will become valueless, causing the Caput to become valueless. Thus, the reduced cost of issuing at lower yield rates offsets the loss on the Caput. If interest rates decline sharply, the firm will gain from issuing at significantly lower yield rates."

Rather than encode thousands of explicit rules for each possible behavior of every possible subset of fundamental factors, we use a causal reasoning engine to derive rules and justifications similar to the ones above whenever needed. This engine uses structural equations stored in a class node in a simulation that generates the behavior of each class of instrument for the event being hedged, and compares this against the trader's protection profile. Since traders most often express both protection profiles and predictions of fundamental factors' behavior in qualitative terms, we use a qualitative causal reasoning engine similar to Kuipers' QSIM [6] system to work off such input.

3.3 A Calculus to Combine Knowledge of One-Dimensional Order

Experts appear to use another type of compiled rule to rank hedge vehicles along various attributes. For example, consider the rule "If fluctuations in the firm's credit spread¹ are small then buy exchange traded options on Treasury rates rather than over the counter options on corporate bond rates."

Experts justify such rules based on knowledge of rankings of vehicles along attributes such as *customization* and *unwind-cost*, and of dependency relationships between attributes such as "the less customized an instrument the cheaper it is to unwind." For example, "An exchange traded option is preferable on an over the counter option since it is less customized and thus less costly to unwind before expiration, and since there is no need to hedge the credit spread."

Experts rank vehicles along qualitative-valued attributes using what [5] refers to as commonsense knowledge of one-dimensional order. Such knowledge is equivalent to the state of knowledge of a total order holding between vehicles along a specific attribute in the real world. It is a partial knowledge of that total order, and thus corresponds to a partial order. For example, a ranking of A, B, C, D along *liquidity* can be $(A > \{B\ C\} > D)$, where $\{B\ C\}$ means it is unknown whether B is more or less liquid than C. The state of knowledge of the order relation of vehicles along quantitative-valued attributes usually corresponds to a total order.

The various attributes along which experts rank vehicles can be organized in a hierarchy based on dependency relationships holding between them. We shall refer to the lowest level attributes in that hierarchy as *atomic attributes*². While rankings of vehicles along atomic attributes can be obtained directly from the knowledge

¹Yields on corporate bonds are often expressed as a spread over the yield of corresponding maturity Treasury securities.

²Experts rank vehicles along over 45 attributes, of which about 30 are atomic.

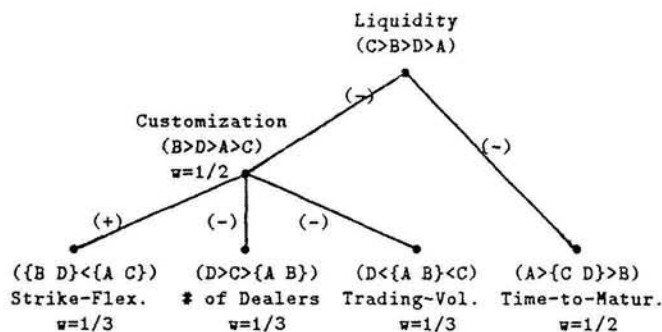


Figure 2: Derivation of rankings by propagation of partial orders

stored in each node in the network, rankings along non-atomic attributes can be derived from rankings along lower level attributes. Thus, comparing two vehicles along a non-atomic attribute requires propagation of the effects of the rankings along lower level attributes for the two vehicles. For example, consider the four vehicles A, B, C, D that need to be ranked along *customization* and *liquidity*. Assume that *customization* is equally determined by the atomic attributes — *strike-flexibility*, *number-of-dealers*, and *trading-volume* — and that *liquidity* is equally determined by the atomic attribute *time-to-maturity* and the non-atomic attribute *customization* (see Figure 2). A propagation of known rankings along atomic attributes up in the hierarchy in Figure 2 determines the rankings along *customization* and *liquidity*.

The ranking of vehicles along non-atomic attributes depends on rankings along atomic attributes and on the weights assigned to them in the propagation of rankings. These weights can change from one case to another. In fact, they are determined based on how the individual trader prioritizes desired characteristics of the hedge vehicle and on the relative importance he assigns to each atomic attribute under the predicted state of capital markets³. Thus, derived rankings of vehicles along non-atomic attributes is also determined by the way the hierarchy of attributes is parameterized.

4 Designing Hedge Vehicles

The way experts apply rules while designing hedge vehicles suggests that hedge design can be seen as a *screening process*. Starting with a set of candidate vehicles, a sequential application of constraints is used to eliminate inferior candidates. In this scenario, constraints are

³A change in the state of capital markets is predicted from the change in the state of fundamental factors (e.g., inflation) and of market indicators (e.g., the Dow-Jones Index).

used as justifications for rejecting hedge candidates. To implement that screening process we have identified the constraints and the knowledge about attributes of instruments that is used to reason about each constraint.

4.1 The Input

The input to the hedge design process is a description about the event to be hedged, the hedging environment, and the hedger. Knowledge about the event includes a qualitative description of the potential future behaviors of fundamental factors associated with the event (i.e., direction of change over a time period), the timing of the event, and the hedger's degree of belief in its occurrence. For example, "interest rates are very likely to increase during the next six months."

Knowledge about the environment includes a qualitative description of the predicted behavior of other relevant fundamental factors over the hedge period.

Knowledge about the hedger includes design constraints of two types — *feasibility constraints* and *optimality constraints*. *Feasibility constraints* always include the trader's protection profile and the hedge period length. Other feasibility constraints are situation dependent and are used to specify policies and preferences of the hedger. These can be monetary constraints such as "use vehicles that require no up-front cash," regulatory constraints such as "use vehicles that maintain the value of the hedger's non-liquid assets under a certain percentage of his capital," position limits such as "do not buy May-91 Treasury bonds if their price exceeds \$100," and miscellaneous constraints such as "do not write options."

Optimality constraints are used to describe desired properties of hedge vehicles. These constraints apply to every situation, but their priorities change from one situation to another. For example, in case of a firm whose plan to issue debt is uncertain, the optimality constraints "maximize liquidity" and "minimize unwind-cost" are likely to be assigned high priorities.

4.2 The Representation in Action

Suppose we are given a set of candidate instruments each represented by an instance node in an ON-ISA network. The idea is to ask every node that is not yet rejected if it can satisfy certain feasibility constraints. A node that cannot satisfy any one constraint, rejects itself and keeps a reference to that constraint as the justification.

The first two feasibility constraints are applied only to class nodes. One constraint enforces the underlying principle of hedging according to which assets with similar or opposing sensitivities must be matched. In effect, every class node whose valuation model is not sensitive to

The situation

The current interest rate (R) is r_t . The hedger believes that interest rate is likely to behave in one of two ways. There is a small possibility that it will remain steady, and a strong possibility that it will fluctuate either way (i.e., move below r_1 or above r_2 , but no higher than r_3). The trader needs a hedge vehicle that will reduce the yield rate he should pay on bonds he plans to issue in the near future, in case his projections turn out to be correct.

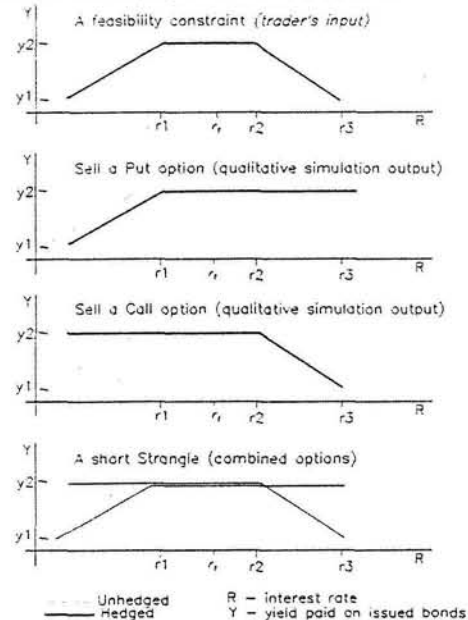


Figure 3: Construction of a compounded vehicle

the fundamental factor associated with the event rejects itself and its instances.

The other feasibility constraint is applied to every non-rejected class node to identify classes of instruments that provide the desired protection profile. The causal reasoning engine is used to do a qualitative simulation for each class node, and the result — a description of the qualitative behavior of instruments in that class under the event to be hedged — is stored in each class node. A qualitative simulation is performed twice for each class, one for a *sell* position and one for a *buy* position (i.e., two vehicles that can be constructed from an instrument). Every class node whose behaviors do not match the desired protection profile rejects itself and its instances. If all candidate nodes are rejected, compounded vehicles (new feasible solutions) are created. That is, vehicles that provide protection profiles that partially match the desired protection profile are used to construct new vehicles that provide the desired protection profile. For example, a compounded vehicle called *short strangle* can be constructed from a "sell Puts" vehicle and a "sell Calls" vehicle (see Figure 3). The newly created compounded vehicles are added as instances of the class node *compounded-vehicles* and marked as candidates.

Additional feasibility constraints can be related to the

amount of up-front cash or the type of collateral required to set up the necessary vehicle. Some of these are applied only to every remaining class node, while others are applied to every remaining instance node.

Last, we apply *optimality constraints* such as “minimize *unwind-complexity*.” The remaining candidates are ranked along each attribute describing an optimality constraint. First, a ranking along every atomic attribute is constructed using knowledge stored in each node about that attribute. Then, the hierarchy of attributes is parameterized. That is, weights for each ranking are derived from the trader input regarding the predicted change in the state of fundamental factors and based on his prioritization of optimality constraints. Finally, rankings along atomic attributes are propagated up in the hierarchy of attributes to derive rankings along non-atomic attributes.

If one candidate dominates all others on all optimality constraints, a best hedge vehicle was found. Otherwise, a small subset of vehicles that rank high on most optimality constraints is selected. In such case, the trader can do two things. One is request the system to construct one vehicle in the subset synthetically. For instance, if one vehicle is an exchange traded future that ranks low only on *maturity-match* due to an early maturity, the system uses two similar futures, one with an earlier and one with a later maturity, such that when combined they create a synthetic future whose maturity matches the end of the hedge period. The other thing the trader can do is find a counter party with whom he can trade over the counter a customized version of one of the vehicles in the subset.

5 Current Status

An appropriate technology for the implementation of an expert system for hedging must primarily provide tools that support the object-oriented nature underlying our representation. After considering various object-oriented tools (e.g., SMALLTALK, FLAVORS, C++) we have decided to use the Zortech C++ package for three reasons. First, it provides several object-oriented capabilities that other tools do not provide (e.g., multiple inheritance). Second, it allows for portability across platforms — something that is important for the integration of our system with other systems dealing with different aspects of risk management. Finally, it supports integration of system modules that were developed using other tools.

So far we have implemented all of the architectural components, the causal reasoning engine in COMMON LISP and the other components in C++. We are currently setting up a small knowledge base to run the model and compare its behavior with that of experts.

The prototype is to be tested to determine the appro-

priateness of the representation and reasoning techniques we are using. We propose to describe several real-world hedging situations to several traders. Each trader will be requested to verbalize his thoughts in the course of deciding what vehicles to use in each situation. Then, for each case, traders’ behavior will be recoded using protocol analysis. The same situations will be given to the system, and the system’s behavior consisting of explanations of decisions it makes while applying different constraints will be recoded. The system’s behavior will be compared to that of traders along three dimensions: (1) the constraints applied in each situation (i.e., search-scope correspondence); (2) the degree of match of generic reasoning activities used with each constraint (i.e., operational correspondence); and (3) the degree of match between domain entities, relationships, attributes, and values that appear in explanations of traders versus those of the system.

Initially, the emphasis will be on the system’s ability to demonstrate reasoning capabilities that traders use in hedge design, rather than on the quality of its recommendations.

There might be two possible reasons for discrepancies between the system’s behavior and that of a trader. One is simply the lack of relevant knowledge, which can be elicited from experts and added to the knowledge base without affecting its usability and without a need to modify the representation. The other reason is inability of the representation to handle other types of knowledge that were not identified prior to the testing.

6 Open Questions

There are a number of problems that the system may encounter in the real-world. One is the potentially explosive number of ways to construct compounded vehicles. It can occur in two situations. One involves multiple constructions of a certain compounded vehicle. For example, in principle a vehicle called *Straddle* can be constructed in two ways — “sell one Put and one Call on stock S” and “sell two Calls on S and buy 100 shares of S” (see Figure 4). In effect the number of ways to construct a *Straddle* can be explosive if one considers every instrument whose sensitivity to the hedged factor is similar to that of S. The other situation involves the construction of non-generic derivatives. For example, in a case where the vehicle “buy a Put” provides the desired protection profile, traders may prefer the derivative vehicle “buy a Call on Put (Caput)” since it requires less up-front cash. Clearly, the number of compounded derivatives can become explosive if one creates derivatives of several higher levels (e.g., “Put on a Caput”, etc.).

So far we have not identified what heuristics good

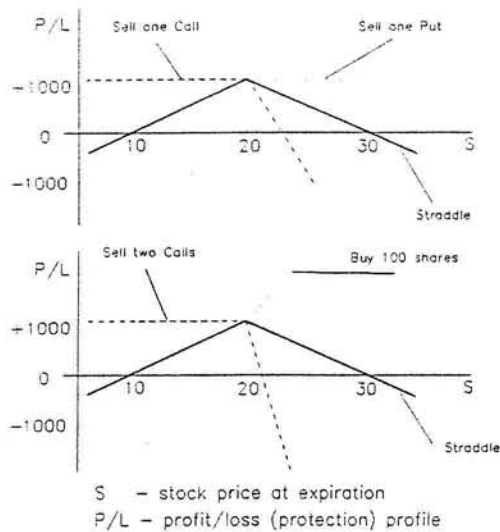


Figure 4: Two ways to construct a *Straddle*

traders use to control the number of ways to construct compounded vehicles. Currently the system is set to create in the first situation one new vehicle for every one of the possible ways, and to create in the second situation only derivatives that are one level higher.

A related problem is the system's inability to construct compounded vehicles which are often referred to as *calendar vehicles*. So far we have discussed only the construction of vehicles that protect against events that conform to traders projections over one particular time horizon. There are cases, however, where traders make short, medium, and long terms projections (e.g., "interest rates will increase in the next month and decrease in the five months that follow"). To benefit from such projections traders can use option-based vehicles that exploit the time decay effects on the value of options of different maturities. That is, the time decay is large on a day-to-day basis in options of under one month's maturity, but is negligible for options of over three months' maturity. For example, the calendar vehicle "sell a near option" and "buy a far option" can be used to benefit from the disproportional change in the value of the two options under specific projected market conditions. We, however, are under the impression that traders use such vehicles very infrequently, mainly due the inherent complexity involved in making projections that conform to multiple time horizons [7].

Another problem is the extensive need for user input for the parameterization of the hierarchy of attributes we use to rank vehicles along optimality constraints. Currently, the weights that specify the importance of every dependency in that hierarchy are derived from the trader's prioritization of optimality constraints and

from input describing the current and predicted state of fundamental factors. We suspect, however, that these weights do not change significantly from one case to another. While considering the trade-off between the payoff from using better hedge vehicles and the costs involved in the design of better vehicles, the costs of gathering and analyzing market information about a large number of economic factors that traders use to make predictions, it is possible that good traders concentrate on a small number of key factors, based on which they make certain assumptions about other factors. We believe that a study of how traders identify a different set of key factors in each case might reveal what of default reasoning scheme can enable the system to derive these weights with less user input.

One other problem involves the relaxation of feasibility constraints. In some cases it is possible that due to a constraint such as "use only vehicles that require no upfront cash" the system will conclude that there are no feasible hedge vehicles. It may be possible, however, that the relaxation of certain feasibility constraints can resolve such situations. Currently, the system cannot determine which constraints must be relaxed in every case to allow it to recommend at least one hedge vehicle. Rather, the user must first determine which constraints to relax, and then rerun the system.

References

- [1] Baecher, E., and Goodman, S.L., *The Goldman Sachs Guide to Hedging Corporate Debt Issuance*, Financial Strategies Group, Goldman Sachs, 1988.
- [2] Dhar, Vasant, and Pople, Harry, "Rule-Based Versus Structure-Base Models for Explaining and Generating Expert Behavior," *Comm. of the ACM*, Vol. 30, No. 6, June, 1987.
- [3] Elton, J.E. and Gruber, J.M., *Modern Portfolio Theory and Investment Analysis (3rd edition)*, John Wiley & Sons, Inc., 1987.
- [4] Figlewski, S., "Topics in Stock Index Futures and Options: Insuring Portfolio Insurance," *Equity Research Notes*, The First Boston Co., Nov. 18, 1987.
- [5] Kuipers, B., "On Representing Commonsense Knowledge," in Findler, editor, *Associative Networks*, 393-408, 1979.
- [6] Kuipers, B., "Qualitative Simulation," *Artificial Intelligence*, 29:289-338, 1986.
- [7] Sutton, William, *Trading in Currency Options*, New York Institute of Finance, 1988.