

1995

Controlling the Complexity of Investment Decisions Using Qualitative Reasoning Techniques

Michel Benaroch

Syracuse University, School of Management

Vasant Dhar

New York University, Stern School of Business

Follow this and additional works at: <http://surface.syr.edu/mgt>



Part of the [Business Administration, Management, and Operations Commons](#)

Recommended Citation

Benaroch, Michel and Dhar, Vasant, "Controlling the Complexity of Investment Decisions Using Qualitative Reasoning Techniques" (1995). *Management*. Paper 6.
<http://surface.syr.edu/mgt/6>

This Article is brought to you for free and open access by the Whitman School of Management at SURFACE. It has been accepted for inclusion in Management by an authorized administrator of SURFACE. For more information, please contact surface@syr.edu.

Controlling the Complexity of Investment Decisions Using Qualitative Reasoning Techniques

Michel Benaroch

Assistant Professor of Information Systems

School of Management, Suit 400

Syracuse University

Syracuse, NY 13244

Phone: (315) 334-3492/1385

E-mail: mbenaroc@mailbox.syr.edu

and

Vasant Dhar

Associate Professor of Information Systems

Stern School of Business

New York University

Management Education Center

44 West 4th Street, 9-170

New York, NY 10012

Abstract

Assembling financial instruments such as equities, bonds, options, and other derivatives into a portfolio requires a thorough understanding of how the portfolio will behave in response to changes of specific economic variables and parameters of the instruments. With more information about a more diverse set of instruments becoming available to traders, it is becoming important to limit the complexity of the analysis involved. We show how this complexity can be limited by using qualitative analysis, where the objective is to construct a few good vehicles which can then be analyzed quantitatively. We illustrate how two qualitative reasoning techniques – *qualitative simulation* and *qualitative synthesis* – are used to design investment vehicles for risk management purposes. These techniques are currently employed by a prototype expert that aims at assisting traders solving a risk management problem called hedging.

Key Words

Qualitative Reasoning (QR), Qualitative Reasoning Techniques, Qualitative Simulation (QSIM), Qualitative Synthesis (QSYN), Investment Decisions, Financial Risk Management, Payoff-Profile, Risk Management Vehicle, Vehicle Configuration, Financial Instruments

1 Introduction

Hedging is a risk management area that is concerned with the design of vehicles which eliminate losses due to, or generate profits from taking, risks that are associated with uncertain events, such as unanticipated changes in currency exchange rates [9]. A risk management vehicle is created by buying and/or selling instruments, such as: bonds, stocks, and options. It allows a trader to control the balance between risk and reward, as a function of how the predicted state of the economy will affect the value of various instruments.

The design of risk management vehicles starts by identifying configurations – combinations of instruments – that provide some goal *payoff-profile* (p-p). The goal p-p specifies what a trader is willing to pay and what s/he is willing to risk to derive a certain payoff pattern. It is defined based on qualitative assessments of how the predicted behavior of certain economic variables (e.g., interest rate, oil prices, demand for orange juice) is likely to affect the value of instruments available in the marketplace. Many of these assessments are based on current market information and/or historical trends.

Example: A firm with a loan that is up for renewal in six months believes that the interest rate will rise from its current level r_t to r_c , with a chance it will decline no lower than r_f . If the interest rate rises (declines), the firm will pay a higher (lower) rate on its loan. The firm therefore defines the ‘cap-floor’ p-p in Figure 1a. This p-p states that the firm seeks to pay loan rate l_c if interest rate rises above r_c , and to take the risk of paying rate l_f if interest rate declines below r_f .

Solution: One vehicle configuration that provides this p-p is created by selling one *call* option on bond B with strike price b_1 and buying one *put* option on B with strike price b_2 , $b_1 < b_2$ (Figure 1b). A *call* (*put*) option on B gives its buyer the right to buy (sell), and obligates its seller to sell (buy), B for the agreed upon strike price at some future expiration date. Hence, if interest rate rises above r_c , the price of B will decline below b_2 , causing the put to become more valuable as the firm can gain from selling B for b_2 . This gain will offset the increase in loan rate, and the actual loan rate will be l_c (Figure 1c). If interest rate declines, the price of B will rise, making the put worthless. But, the reduced loan rate will offset, and more, the put’s purchase cost. At the same time, if interest rate is above r_f , the firm will profit from the cost received for the call sold. If it declines below r_f , the call sold will

be exercised by its buyer, causing a loss to the firm. However, this loss will be offset by the lower loan rate, and the actual loan rate will be l_f (Figure 1c).

INSERT FIGURE 1 ABOUT HERE

This example shows that a vehicle configuration is specified in terms of: the combination of instruments to be purchased and/or sold, the ordinal relation of these instruments' strike prices, and the unit proportions to be purchased/sold of each instrument. Moreover, in order to configure vehicles that provide some goal p-p, one needs knowledge of the profit&loss pattern (or p-p) provided by each of the instruments traded in capital markets, and of how to 'synthesize' the goal p-p by permuting the p-p's of specific instruments.

The problem of configuring vehicles can be characterized by three features. One, the number of alternative vehicle configurations is combinatorial, given the thousands of individual instruments available to the trader. Second, the alternative configurations change over time with the constant issuance of new instruments and the elimination of matured ones. Last, each instrument may provide different p-p's under different market situations, depending on its sensitivity to the one or more particular economic variables being hedged.

The difficulty that a typical trader faces while configuring vehicles is mainly due to her/his specialization in only few types of instrument. For example, a trader who specializes in stocks of the automobile industry may not know enough about how to manage risk using the variety of other instruments (e.g., bond options, Eurodollar futures). Without computerized support, s/he is prone to consider only a subset of the alternative vehicle configurations for a given situation, and thus to make suboptimal decisions.

In considering a computer-based solution for supporting traders in the configuration of vehicles (and the rest of the hedge design process) it became apparent to us that the search space of alternative configurations is very large and constantly changing. Hence, the problem cannot be solved using conventional optimization

techniques, such as linear programming. Also, pre-storing all possible configurations for selection is not feasible. Thus the configurations must be constructed. However, given the explosive number of usable instruments, constructing configurations is a combinatorial generate-and-test search problem. It is therefore important to use good heuristics to constrain the generator.

The earlier example shows that one can use two means (or heuristics) for this purpose. One is ‘qualitative abstraction’, which allows reasoning about entire classes of similar instruments (e.g., all call options on bond B with strike price b_i) instead of individual instruments (e.g., a June-93 call on a March-95 T-bond with \$45 strike price). This allows replacing the p-p’s of all individual instruments of the same type by one qualitative p-p, thus reducing significantly the number of p-p’s that has to be permuted in order to synthesize the goal p-p. However, such an abstraction causes lose of useful information. As Figure 1b illustrates, to find the right ordering of strike prices, the first element in the qualitative p-p of “buy put” had to be made longer than that in the p-p of “sell call”. Hence, to rediscover some of the information lost, one ought to know how to manipulate the qualitative p-p’s used to synthesize the goal p-p. The other means is the use of a qualitative causal calculus. Specifically, one can predict the situation-specific p-p of an instrument by qualitatively analyzing causal relationships between economic variables and the value of that instrument (i.e., “If interest rate rises, bond prices decline, causing the value of bond puts to ...”). This eliminates the need to pre-store the various p-p’s that every instrument can provide under different market situations.

What techniques that employ similar means can be used to configure risk management vehicles? Qualitative reasoning (QR) techniques, which were originally developed to support the analysis and design of physical systems, seem a natural choice. These techniques basically emulate the ability of humans to reason about physical systems using a qualitative causal calculus [8]. More importantly, QR techniques seem suitable because a vehicle can be conceptually analogized to a physical system. Instruments are elementary components which are ‘connected’ in a certain way to provide some desired functionality (i.e., goal p-p). The behavior (i.e., value)

of a vehicle results from the ‘combined’ behaviors of its instrument components, where the behavior of instruments is characterized by the contingent behavior of economic variables.

This paper explains how QR techniques are used to solve the problem of configuring risk management vehicles. The paper is organized as follows. Section 2 reviews the role of QR techniques in the analysis and the design of physical systems, and how this role relates to the problem at hand. Sections 3 and 4 explain and illustrate how two QR techniques – *qualitative simulation* and *qualitative synthesis* – are used to configure risk management vehicles. Section 5 shows how the use of these two techniques fits into the overall problem of designing risk management vehicles.

2 Qualitative Reasoning Techniques

Much of the work on qualitative reasoning (QR) about physical systems (e.g., [4]) relies on the relationships between: *structure* (or *configuration*) – a collection of components connected as a system; *behavior* – a sequence of states that a system and its components exhibit over some time-interval; and *function* – the purpose of structure in producing the behavior of a system. The behavior of a system results from interactions between the behaviors of its components. The effects of a change in the state of one component propagate locally through structural connections causing a change in the state of other components and of the system as a whole. On the other hand, the function of a system explains in terms of causality why and how the structure of a system determines its behavior [6].

One can distinguish between QR techniques that are used to analyze a system, and ones that are used to construct a system. The goal of QR analysis techniques is to infer behavior from structure. For example, *qualitative simulation* [11] accepts two inputs: a model describing the structure of a system in terms of the parameters characterizing that system as well as structural connections between them, and the initial state of parameters in that model. It simulates how changes in the state

of parameters propagate through structural connections in order to predict the qualitative transitions that a system will make over time, and to explain in causal terms how the behavior of that system results from its structure.

The primary goal of QR-based design techniques is to infer the structure of a system from its function. For example, consider a QR-based technique called *interaction-based invention* (Ibis) [13]. The input to Ibis includes knowledge describing each type of elementary component in some domain in terms of the parameters characterizing it (e.g., $h(V)$ for height of fluid in a vat component V) and the types of components with which it can interact (i.e., be connected). It also includes some desired interaction. For example, the interaction $[h(V) - h(B)] = [dh(B)/dt]$, where B is a bowl component and $[]$ denotes the sign of an algebraic expression, states that the goal is to design a device in which the change in the height of fluid in B is a function of the difference between the heights of fluid in V and B . Ibis maps the desired interaction onto all possible chains of interactions between the various types of components, such that each chain relates the parameters in the desired interaction (e.g., $h(V)$, $h(B)$, $dh(B)/dt$). It selects the shortest chains, each as a candidate design. As an interaction is a qualitative (or quantitative) relationship between multiple parameters, Ibis then uses a QR analysis technique to test whether or not a candidate design generates the desired interaction. Apparently, Ibis uses a generate-and-test approach, which can be inefficient when the number of types of elementary components is large.

For some design situations, one can specify function in terms of desired behavior. For example, “cap loan rate”, which is a function of the portfolio, can be specified in terms of the behavior of loan rate in response to interest rate. This kind of behavior can be expressed as a qualitative two-dimensional piecewise linear function, which specifies the input ‘values’ that a system can accept and the corresponding output ‘values’ it should produce. Specifically, if the behavior of each type of elementary components in a domain can be expressed as a qualitative two-dimensional piecewise linear function, the idea is to algebraically ‘compose’ the goal behavior from the behaviors of components. This approach basically identifies components that can

be connected so that the overall behavior resulting from interactions between their behaviors according to laws of causality is identical to the desired behavior. One QR-based technique which applies this approach is called *qualitative synthesis* [2]. Clearly, deriving structure from behavior is a synthesis (not a selection) problem.

The next two sections explain how QR techniques are used in actuality for the analysis and the configuration of risk management vehicles. Hereafter, we shall distinguish between two types of vehicles – *generic* and *compound*. A *generic* vehicle is created by either selling or buying instruments of a single type, whereas a *compound* vehicle is created by selling and/or buying instruments of multiple types. Given that in some situations both a generic and a compound vehicle can provide the same p-p (e.g., the p-p of “buy put on stock S ” is identical to the p-p of “buy stock S ” combined with “sell call on stock S ”), traders usually prefer to first look for configurations of generic vehicles that provide the goal p-p. A generic vehicle is simply easier to set up, maintain, track over time, etc.

3 Configuring Generic Vehicles

To understand how generic vehicles can be configured, consider the following example.

Example: A firm that plans to issues bonds in order to raise capital believes that the risk-free interest rate is likely to increase, with a chance it will decline, prior to the issuance date. As an increase in the interest rate will higher the yield rate offered on issued bonds, the firm wants to protect itself against an increase, while preserving the ability to benefit from a decline in the interest rate. It therefore defines the “cap” p-p in Figure 2.

Solution: One generic vehicle configuration that provides this p-p is created by the purchase of put options on some bond B with strike price b_i . An increase in interest rate will cause the price of B to decline below b_i , allowing the firm to profit from selling bonds for b_i and to offset the extra cost of issuing bonds at a higher yield rate. Alternately, a decline in interest rate will make the put valueless, but allow the firm to issue bonds at a lower yield rate and make a profit to offset, and more, the cost paid for the put.

This example shows that the p-p of a generic vehicle can be derived by qualitatively analyzing causal relationships between the behavior of the economic variable(s) being hedged (e.g., interest rate) and the value of the instruments purchased/sold. In effect, this analysis simulates qualitatively the behavior (i.e., value) of the trader's hedged position, as a function of the behavior (i.e., value) of the purchased/sold instruments (i.e., vehicle) and the contingent behavior of the variable(s) being hedged. *Qualitative simulation* is one technique that can emulate this kind of analysis.

 INSERT FIGURE 2 ABOUT HERE

3.1 Qualitative Simulation

Qualitative simulation (QSIM) is a QR technique that can derive the behavior of a system based on that system's structure [11]. The main ideas behind QSIM are: (1) the structure of a system can be described by equations modelling structural connections between the parameters (i.e., continuous functions) characterizing that system; (2) a change in the state of one parameter propagates to other parameters through structural connections; (3) the qualitative behavior of a parameter can be described by the transitions it makes from one state to another (e.g., a change from an 'increasing' state to a 'steady' state); and (4) the qualitative behavior of a system can be described by the qualitative behavior of every parameter characterizing that system.

QSIM receives as input a qualitative structural model (i.e., a set of qualitative equations) of a system and the initial state of parameters in that model. The qualitative state of a parameter is represented by a pair $\langle qdir, qual \rangle$, where $qdir$ is the qualitative direction of change of that parameter's value (i.e., $qdir \in \{-1, 0, 1\}$ or {decreasing, steady, increasing}) over $qual$ – a qualitative point or region on the real-line. Assuming that a system is in equilibrium and that one or more of its parameters are perturbed, QSIM propagates the effects of the perturbation to other parameters through structural equations according to various calculus laws. For example, given a system whose structure is described by the equation $X = Y + Z$,

if Y is perturbed to start increasing, propagation of this perturbation using limit analysis will conclude that X also starts increasing. By propagating changes in the current state of parameters, QSIM derives the next qualitative state of every parameter in the system and of the system as a whole. QSIM continues to propagate changes in the state of each parameter, until all parameters reach a ‘steady’ state or a boundary *qual*. For example, in the above system, unless some ‘external’ parameter causes X , Y or Z to change its behavior, the system will remain in that state forever. Of course, when the structure of a system is described by a number of equations, it is the various interactions between the behavior of parameters in that system which generate complex behaviors.

3.2 Applying QSIM in Hedging

When the ‘structure’ of the trader’s hedged position is known, QSIM can be used to predict its behavior (i.e., p-p) under the market situation of concern. If the predicted p-p matches the goal p-p, the vehicle configuration used to construct the hedged position is suitable.

As a generic vehicle is conceptually a one component (i.e., instrument) system, the structure of a position that is hedged using a generic vehicle can be described by two things. One is the equation $VHP = VUP \pm VI$, which states that the Value of the Hedged Position is the Value of the Unhedged Position plus (minus) the Value of the Instrument sold (purchased) (i.e., value of generic vehicle used). The other thing is the *valuation model* of the instrument purchased/sold. In Finance, causal relationships between economic variables and the value of a specific instrument are each modeled formally by an equation that specifies how a certain economic variable affects the value of that instrument [7]. The set of equations modelling the major relationships for a specific instrument is called the *valuation model* of that instrument. This model’s analytic solution is typically used to compute the fair market value of that instrument. Since each type of instrument is sensitive to a different set of economic variables, different types of instruments have different valuation models.

To illustrate how QSIM predicts the p-p of a position that is hedged by a particular generic vehicle, consider the example of using a “purchase put option on bond” vehicle to “cap” the cost of issuance (see Figure 2). A p-p is expressed symbolically as a sequence of pairs $((V \langle qdir \ qval \rangle)(VHP \langle qdir \ qval \rangle))$, where V is the variable being hedged, and VHP is the value of the hedged position. The “cap” p-p is thus expressed as:

$$[((R \langle inc \ (0, r_c) \rangle) (HIC \langle inc \ (0, y_c) \rangle)) \\ ((R \langle inc \ (r_c, \infty) \rangle) (HIC \langle std \ [y_c] \rangle))],$$

where R is interest rate, HIC is the hedged issuance cost, and r_c is the interest rate level corresponding to the cap level y_c on the yield rate offered on issued bonds.

The input for QSIM includes the set of qualitative structural equations in Figure 3, and the initial state of parameters in these equations. In principle, to derive the complete p-p of the analyzed position, one can describe the initial state of every parameter for the current state of R and run QSIM twice – for R increasing, and for R declining from its current level. However, we describe the state of parameters when R is zero, and run QSIM only once, letting R increase over the $qval$ range $(0, \infty)$.

 INSERT FIGURE 3 ABOUT HERE

A trace of the states QSIM derives is presented in Figure 4. In the initial state, interest rate is zero, the price of a yield bearing bond is positive (infinite in the limit), the value of a put on that bond is zero, and the issuance cost (hedged and unhedged) is zero because theoretically a firm can offer an infinitesimal yield rate to get investors to buy its bonds. Starting with this state, state 0, the transitions QSIM predicts can be summarized as follows. In state 1, R 's increase causes B to start declining and UIC to start increasing, in compliance with equations 4 and 2, respectively. Since B has not yet reached x , the put's strike price, P remains zero complying with equation 3, and HIC starts increasing to comply with UIC 's increase in equation 1. In state 3, as R continues to increase, B declines below x , and P begins to increase in compliance with equation 3. In turn, HIC becomes steady at

the “cap” level y_c because QSIM ‘assumes’ that the increase in P balances off UIC ’s increase in equation 1. This assumption is based on the notion that (by definition) a hedge vehicle is constructed to balance off changes in the value of the trader’s unhedged position. This is the whole idea behind using the *hedge ratio* to compute the precise number of units to be purchased/sold of the instruments involved [9].

 INSERT FIGURE 4 ABOUT HERE

The p-p of the hedge position being analyzed is embedded in the sequence of states QSIM derives. It is comprised of the states of R and HIC over a *qual* range, as opposed to at a point. These states are enclosed in dashed boxes in Figure 4. A comparison of this derived p-p with the goal “cap” p-p will thus conclude that a “purchase put option on bond” vehicle can be used to cap the cost of issuance.

3.3 Pragmatic Considerations

In order to identify all generic vehicles providing the goal p-p under the particular market situation being hedged, it is necessary to apply QSIM for every individual instrument one can use to create generic vehicles. Though QSIM is effective in producing the p-p of any instrument under any market situation, its use can be associated with extensive amounts of computation [11]. In hedging, this can be a serious inhibiting factor because: (1) there are too many instruments for which QSIM needs to be applied; and (2) each instrument requires two QSIM runs – for a ‘sell’, and for a ‘buy’ action. We therefore apply several means to keep the use of QSIM tractable.

One means is qualitative abstraction (or inheritance) that is based on the domain’s deep structure. As all instruments of the same class (e.g., put options on a Treasury bond) have the same valuation model, QSIM is applied collectively for all instruments of the same class. This concept can be further exploited, given that instrument classes can be organized in an ISA hierarchy, such as the one in Figure 5, based on specialization relationships between them (e.g., Treasury bonds, bills, and notes are all fixed-income instruments). In such an ISA hierarchy, the (qualitative)

valuation model of one class of instruments can be a specialization of the valuation model of another class. For example, the valuation model of bond options is a specialization of the Black-Scholes model, which is used to derive the valuation model of various option types [9]. Accordingly, QSIM is applied only for each instruments class whose valuation model is a generalization of the valuation models of other classes of instruments. Of course, this kind of qualitative abstraction requires reliance on the specific representation used to capture knowledge about instruments (see [1,3] for details).

 INSERT FIGURE 5 ABOUT HERE

Other means for making the use of QSIM computationally tractable involve the application of several heuristics. First, as the sale/purchase of an instrument that is insensitive to the economic variables being hedged is meaningless from a hedging stand point, QSIM is applied only for classes of instruments whose valuation model references the variables being hedged. Second, as the p-p's for a 'buy' and a 'sell' action are symmetrical because trading is a zero-sum game, QSIM is applied only for a 'buy' action and the p-p for a 'sell' action is derived easily by finding the symmetrical p-p of the one QSIM predicts (i.e., the *qdir* in every state of *VHP* is changed from 1 to -1 and vice versa, and the *qval*'s are adjusted accordingly). Last, to eliminate almost completely QSIM's tendency to branch (and sometime explode) due to its reasoning with qualitative values, we apply various domain-specific assumptions regarding values for which the affects of competing tendencies are balanced (see example above), and use mixed qualitative/quantitative values to describe the magnitude of many economic variables appearing in the models being simulated [10].

4 Configuring Compound Vehicles

Traders configure compound vehicles when the goal p-p is more complex than the p-p's generic vehicles provide, or when all the configurations of generic vehicles that QSIM identifies violate other design specifications (e.g., maturity date, amount of

cash available to acquire the vehicle). To understand how compound vehicles can be configured, consider the following example.

Example: A trader who speculates that the price of stock S will increase above s_1 , but not above s_2 , defines the “ratio-spread” p-p presented in Figure 6.

Solution: One compound vehicle configuration that provides this p-p involves the purchase of one call option on S with strike price s_1 and the sale of two call options on S with strike price s_2 , where $s_1 < s_2$. In case of a price movement above s_1 , the purchased call option allows to profit from buying stocks for s_1 to offset, and more, the cost of this call. At the same time, as long as the price of S is above s_2 , the call options sold allow to profit from the cost paid by another party who believes that the price of S will move above s_2 .

This example shows that a compound vehicle is in fact a combination of two or more individual generic vehicles, and that subsequently one can configure compound vehicles by synthesizing the goal p-p as a linear combination of the situation-specific p-p’s of various generic vehicles.

 INSERT FIGURE 6 ABOUT HERE

Synthesizing p-p permutations which match some goal p-p is a combinatorial generate-and-test search problem. Considering only option-based generic vehicles, for example, the number of possible permutations of p-p’s is 2^{4n} , where 4 stands for the p-p’s of “sell call”, “buy call”, “sell put”, “buy put”, and n is the number of different strike prices (n is in the thousands, considering all traded options on different underlying instruments, such as Treasuries, Eurodollar, Futures, and Stocks). However, since the goal p-p is specified qualitatively, one can suggest making the problem tractable by using qualitative abstraction over the p-p’s of all individual generic vehicles of the same type. Specifically, all individual p-p’s with a ‘similar’ shape can be replaced by one qualitative p-p. For example, consider all individual generic vehicles of the type “buy m calls with strike price s_i ”, where s_i and m are different across vehicles. The p-p of each of these individual vehicles has the following shape: the first element has slope 0 over the range $(0, s_i)$, and the

second element has slope m over the range (s_i, ∞) . All these individual p-p’s can be replaced by one qualitative p-p in which the first element is flat over the *qual* $(0, s)$ with $s \in (0, \infty)$ being an arbitrary qualitative strike price, and the second element’s *qdir* is ‘increasing’ with slope 1. Eventually, the use of such a qualitative abstraction over the situation-specific p-p’s QSIM derives leaves us with a small number of what we will hereafter refer to as *generic p-p’s* (see Figure 7).

 INSERT FIGURE 7 ABOUT HERE

While such qualitative abstraction will significantly reduce the complexity of the synthesis problem, it will also result with loss of important information. For instance, consider the example in Figure 8, which describes a compound vehicle configuration that involves two generic vehicles of the same type (i.e., “buy one call”). Since the precise p-p of these two individual vehicles is now represented by the same generic p-p, the goal p-p cannot be synthesized unless the lost information is rediscovered by stretching and/or steepening elements of the generic p-p’s used in a permutation. Hence, if one is to rely on qualitative abstraction, one must also use good heuristics to uncover the information lost.

 INSERT FIGURE 8 ABOUT HERE

4.1 Qualitative Synthesis

Qualitative synthesis (QSYN) is a QR-based technique that solves the above synthesis problem in the domain of two-terminal systems, i.e., systems with one input node and one output node [2]. QSYN relies on the following principles (see Figure 9 and [12]): (1) a system is comprised of elementary components which are connected in series and/or in parallel; (2) the behavior of a system and of each elementary component can be described by a *transfer function* – the ratio of the Laplace transform of the output to the input, with all initial conditions neglected; (3) given the transfer functions of any two elementary components, their algebraic sum (product) is the transfer function of a system made from the two components

connected in parallel (series); and (4) to configure a prospective system, one can apply algebraic operations on the transfer functions of elementary components to create permutations that match the desired transfer function of that prospective system.

 INSERT FIGURE 9 ABOUT HERE

Instead of applying these principles on transfer functions, QSYN applies them on two-dimensional qualitative piecewise linear functions (hereafter, *q-function*), each of which describes the qualitative behavior of one type of elementary components over their operational regions. Assuming that the desired behavior of some prospective system is represented as a q-function and that the behavior of each type of elementary components is also expressed as a q-function (hereafter, generic q-function), the configuration problem can be solved by algebraically creating permutations of generic q-functions that match the objective q-function. Since in risk management the p-p of a compound vehicle is an additive linear combination of generic p-p's, we shall explain how QSYN solves the problem for the case of elementary components which are connected in parallel.

QSYN receives as input a goal q-function, G , and a set of generic q-functions \mathcal{Q} . It creates one permutation of q-functions at a time as the sum of two different q-functions in \mathcal{Q} . A newly created permutation is a q-function that is then compared against G . If it matches part, or all, of G , that q-function is added to \mathcal{Q} with a reference to the two q-functions in \mathcal{Q} that create it. QSYN repeats the same operations for every pair of different q-functions in \mathcal{Q} , including ones containing q-functions newly added to \mathcal{Q} . In so doing, QSYN finds all permutations of generic q-functions that match G .

Since the number of possible permutations of generic q-functions in \mathcal{Q} can be large (i.e., $O(n^2)$ – n is the number of q-functions in \mathcal{Q}), QSYN constrains its generator by applying knowledge about the additivity of qualitative behaviors (using the transition rules QSIM employs). For example, if the *qdir* of the first element in two q-functions is ‘increasing’, their sum will not match a goal q-function whose first

element $qdir$'s is 'steady'. Also, QSYN applies two heuristic synthesis operators – *STRETCH* and *STEEPEN* – on elements of the q -functions in a permutation. These are used to rediscover the information lost by qualitatively abstracting the behaviors of all generic components of the same type, as illustrated in Figure 8. Before we can demonstrate how QSYN operates, it is necessary to define the sum of two q -functions and the conditions under which two q -functions match.

The sum of two q -functions, denoted \oplus , is defined as follows. A q -function is a sequence of elements, each of the form $((IN \langle qdir \ qval \rangle)(OUT \langle qdir \ qval \rangle))$. Assume the existence of q -functions Q_1 and Q_2 , with m and n elements respectively, and let $[]$ denote the k -th element of a q -function. Elements $Q_1[i]$ ($1 \leq i \leq m$) and $Q_2[j]$ ($1 \leq j \leq n$) are said to be *corresponding*, if the *IN-qval* of $Q_1[i]$ is contained in the *IN-qval* of $Q_2[j]$, or vice versa. We define the *sum of two q -function elements*, denoted $Q_1[i] \oplus Q_2[j]$, to be a new element, $Q_3[k]$, in which: (1) the *IN-qval* is the intersection of *IN-qval*'s of $Q_1[i]$ and $Q_2[j]$; and (2) the *OUT-qdir* is the algebraic sum of *OUT-qdir*'s of $Q_1[i]$ and $Q_2[j]$. The fact that this definition disregards the *IN-qdir*'s and the *OUT-qval*'s does not limit the generality of the synthesis approach used by QSYN. The *IN-qdir*'s are ignored because they are always 1 in every element of any q -function, and the *OUT-qval*'s are disregarded because traders look at their value only during the quantitative analysis and refinement of vehicle configurations. The next example illustrates how the sum of two q -function elements is computed based on this definition, assuming that $(i_1, i_2) \subseteq (i_1, i_3)$:

$$\begin{array}{rcl}
 Q_1[i] & = & ((IN \langle * (i_1, i_2) \rangle) (OUT \langle 1 * \rangle)) \\
 Q_2[j] & = & ((IN \langle * (i_1, i_3) \rangle) (OUT \langle -1 * \rangle)) \\
 \hline
 Q_1[i] \oplus Q_2[j] & = & ((IN \langle * (i_1, i_2) \rangle) (OUT \langle 0 * \rangle)).
 \end{array}$$

Following the above definition we define the *sum of two q -functions*, denoted $Q_1 \oplus Q_2$, to be the sum of every pair of corresponding elements in Q_1 and Q_2 .

Two corresponding elements are matching, denoted $Q_1[i] \equiv Q_2[j]$, if they have the same *OUT-qdir*. For example, although the above two sample elements are corresponding because $(i_1, i_2) \subseteq (i_1, i_3)$, they do not match because the *OUT-qdir*

of $Q_1[i]$ is 1 whereas the *OUT-qdir* of $Q_2[j]$ is -1 . Two *q-functions* are *matching*, denoted $Q_1 \equiv Q_2$, if each pair of corresponding elements in Q_1 and Q_2 match. A *q-function* Q_1 *partially matches* another *q-function* Q_2 , if Q_1 matches the first few consecutive elements of Q_2 .

Now that we know what is the sum of two *q-functions* and what are the conditions under which two *q-functions* match, let us use an example to illustrate how QSYN uses operators *STRETCH* and *STEEPEN*.

4.2 Qualitative Synthesis in Hedging: An Example

Suppose we are trying to synthesize the “ratio-spread” p-p denoted RS in Figure 6. One of the permutations of p-p’s QSYN tries includes the pair of generic p-p’s denoted P_i and P_j in Figure 7. Apparent from the example in Figure 6, RS can be synthesized from P_i and P_j . However, $P_i \oplus P_j \not\equiv RS$ because P_i and P_j are each an abstraction of an entire ‘class’ of individual p-p’s with the same qualitative shape. QSYN therefore tries to use operators *STRETCH* and *STEEPEN* in order to synthesize RS using these two generic p-p’s.

Figure 10 traces QSYN’s synthesis of RS from P_i and P_j . Starting with the first triplet of elements, QSYN concludes that $P_i[1] \oplus P_j[1] \equiv RS[1]$ (Figure 10a). Proceeding with the next triplet QSYN concludes that $P_i[2] \oplus P_j[2] \not\equiv RS[2]$, because the *VHP-qdir* of $RS[2]$ is not equal to the *VHP-qdir* of $P_i[2] \oplus P_j[2]$ (Figure 10b). However, since the *VHP-qdir* of $RS[2]$ is equal to the *VHP-qdir* of $P_i[2] \oplus P_j[1]$, a modified version of P_j (denoted P_j' in Figure 10b), in which the first element is stretched over the *V-qual* $(0, s_2)$, is more likely to contribute to the synthesis of RS . QSYN therefore uses operator *STRETCH* to extend $P_j[1]$ over the range $(0, s_2)$, and to conclude that $P_i[2] \oplus P_j'[1] \equiv RS[2]$. For the next triplet of elements QSYN concludes that $P_i[2] \oplus P_j'[2] \not\equiv RS[3]$, because the *VHP-qdir* of $RS[3]$ is not equal to the *VHP-qdir* of $P_i[2] \oplus P_j'[2]$. However, this mismatch can be eliminated by modifying the *VHP-qdir* of $P_j'[2]$ from -1 to -2 . QSYN therefore applies operator *STEEPEN* to create a new version of P_j' (denoted P_j'' in Figure 10c), and to conclude that $P_i[2] \oplus P_j''[2] \equiv RS[3]$. At this point QSYN found a full match.

 INSERT FIGURE 10 ABOUT HERE

The matching permutation QSYN has synthesized is made from two generic p-p's which were modified by operators *STRETCH* and *STEEPEN*. These modified p-p's provide important information about how to configure “ratio-spread” vehicles. First, P_j and P_j'' have the same qualitative shape of the p-p's of a “buy call option on some stock S ” vehicle and a “sell call option on some stock S ” vehicle, respectively. Second, since $s_1 \leq s_2$, the strike price of the purchased call (s_1) should be smaller than that of the sold call (s_2). Last, the absolute value of the *VHP-qdir* of the second element in P_j'' is 2, something which indicates the need to sell more than one call for every call purchased. This information is identical to the one provided by the example in the beginning of Section 4.

4.3 Computational Feasibility

Though the number of p-p permutations QSYN has to analyze is explosive, the use of operators *STRETCH* and *STEEPEN* allows QSYN to avoid searching exhaustively the space of all permutations of generic p-p's. When these two operators are applied under the right conditions (see [2] for details), they can narrow down significantly the search space. Figure 11, for example, shows the search tree only for one permutation involving p-p's P_i and P_j in the synthesis of a “ratio-spread” p-p. The arrowed branches in the tree are the ones QSYN explores, while all other branches are readily pruned.

 INSERT FIGURE 11 ABOUT HERE

One can probably improve the efficiency of QSYN by finding all the specific conditions under which operators *STRETCH* and *STEEPEN* are more likely to lead to a successful synthesis of a goal p-p. Nevertheless, our experience with a C++ implementation of QSYN running on a 386-based PC indicates that QSYN's performance is adequate for its intended application. For example, in the case of a goal p-p with six linear elements and an input of seven generic p-p's, QSYN takes

only a fraction of a second to produce all the possible permutations matching that goal p-p.

5 The Other Parts of Hedge Design

Hedge design can be viewed as a constrained multi-objective optimization problem [1,3]. Constraints exist for matching maturity dates, not exceeding available resources in setting up the hedge vehicle, and so on. The multiple objectives are maximizing liquidity, maximizing maturity match, minimizing setup cost, and minimizing credit risk, among others. Solving such a problem is hard. Performing qualitative simulation/synthesis is one simplification strategy, but it focuses only on matching the goal payoff-profile, ignoring some of the above constraints and objectives. Upon applying the configuration results of qualitative simulation/synthesis for the construction of individual vehicles (using the many instruments traded in capital markets) one usually ends up with a large number of candidate vehicles. These candidate vehicles need to be screened down to a manageable number by applying design constraints, and then ordered using the optimization criteria. At that point, quantitative analysis can be performed to determine expected values of payoffs under varying parameters of the instruments involved, their price volatility, their time to maturity, and so forth.

The screening part of the task is fairly straightforward. Individual candidate vehicles are analyzed against design constraints, such as: the cash upfront fee needed to set up a vehicle must not exceed the amount of cash available to the trader, the tax regulations associated with a vehicle must grant the tax benefits sought by the trader, and the minimum size contract of a vehicle should not exceed the value of the asset being hedged. For certain feasibility constraints, the ISA hierarchy of instrument classes (Figure 5) permits using qualitative abstraction and/or inheritance in order to rule out entire classes of instruments. For example, if the hedger does not want to use over the counter instruments, ruling out a general class of over the counter instruments (e.g., call on put on bonds) will automatically rule out all the subclasses and instances of that class.

Ordering the remaining vehicles must necessarily be done by the trader, since none of the remaining solutions is Pareto-Optimal to others. In other words, trade-offs must be explored actively by the trader, based on personal preferences and risk taking propensity. Vehicles can be ordered based on the importance assigned to the various objective functions, which include: liquidity (maximize), setup cost (minimize), unwinding complexity (minimize), and maintenance complexity (minimize), among others. Specifically, these ‘competing’ objectives are prioritized based on preferences of the trader, as well as qualitative assessments of the anticipated market conditions (i.e., state of various economic variables) at the end of the hedging period and their effect on the behavior of vehicles. For instance, consider the ‘maximize liquidity’ objective, which may initially be given a high preference. Yet, this objective may have an overall low priority for vehicles involving put options on bonds, if the trading volumes of put options on bonds are expected to be high (for example, because trade balances and subsequently foreign investment are expected to be high).

Finally, a quantitative what-if analysis is required to see the ‘precise’ protection level provided by each candidate vehicle. This analysis helps evaluate tradeoffs between candidate vehicles in response to changes in economic variables or instrument parameters. To illustrate one of these tradeoffs, consider the following hypothetical situation. A trader who speculates that the value of stock S will increase in the near future has to select among two candidate vehicles that involve the purchase of a different call option on S . The primary parameters that determine the attractiveness of, or the precise profit&loss (P&L) pattern expected from such an option are: c – the purchase price of the option, s – the current price of stock S , σ – the currently observed price volatility of stock S , x – the option’s strike price, t – the option’s time to maturity, and r – the risk-free interest rate [5]. The P&L patterns offered by two call options on S , C_1 and C_2 , where C_1 has a lower price and an earlier maturity date, are presented in Figure 12 (two tangent curves). Evidently, C_1 is preferable because it offers a P&L pattern that is slightly more attractive, *assuming that the price volatility of S will not change over the hedging*

period. However, C_2 may appear attractive in another respect, namely: its value is more sensitive to changes in the price volatility of S (i.e., $\partial c_1/\partial\sigma < \partial c_2/\partial\sigma$). The latter observation means that, if the price volatility of S will rise during the hedging period, C_2 will end up offering a more attractive P&L pattern (see Figure 12, top curves). Hence, if the trader believes that the price volatility of S is going to increase, s/he will prefer C_2 over C_1 , even though C_2 is more costly. Of course, by preferring C_2 the trader indicates that s/he is willing to take the risk that the price volatility of S will decline, in which case C_2 will end up offering a much less attractive P&L pattern (see Figure 12, bottom curves). This example shows that tradeoffs between candidate vehicles are typically assessed based upon the trader's personal beliefs and risk attitude.

 INSERT FIGURE 12 ABOUT HERE

6 Concluding Remarks

Current practice in organizations indicates that traders have narrow areas of specialization and expertise. By focusing only on a limited number of types of instruments (foreign exchange, muni bonds, etc), traders reduce the complexity associated with putting together hedge vehicles to a manageable level. At the same time, however, traders risk making suboptimal decisions in that they overlook a whole range of vehicles, some of which may occasionally be most suitable for their risk management needs.

The trend in financial institutions, however, is towards integration, where more global information is becoming available to traders. Under this scenario, traders will begin to seek out better solutions, as long as the right kinds of tools are available to help them manage the additional complexity brought about by consideration of a larger set of instruments. For the most part, this involves doing a lot of screening for the trader, and presenting only the most promising vehicles for quantitative analysis.

Our objective is to show how we can capitalize on the integration. By considering a more diverse set of instruments, better risk management vehicles can be constructed. In this paper, we have presented a model that deals with the additional complexity brought about by considering the extremely large set of instruments available in capital markets. This model has been implemented in a prototype expert system for hedging [1,3]. So far, our experience with this system suggests that the model we have presented performs well its intended function.

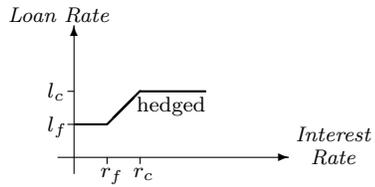
References

- [1] M., Benaroch, *On Building Deep Expert Systems for Wide and Volatile Business Domains*, Unpublished Doctoral Dissertation, Stern School of Business, New York University, 1992.
- [2] M., Benaroch, and V., Dhar, "Qualitative Synthesis of System Configurations Based on Desired Behavior," *Proceedings of the 9th Canadian Conference on Artificial Intelligence*, Vancouver, CANADA, IEEE Press, 1992.
- [3] M., Benaroch, and V., Dhar, "A Knowledge-Based Approach to Solving Hedge Design Problems," *Proceedings of the 1st International Conference on AI Application on Wall Street*, New York, IEEE Press, October, 1991.
- [4] D., Bobrow, *Qualitative Reasoning about Physical Systems*, Elsevier North-Holland, New York, NY, 1984.
- [5] J.C., Cox, and M., Rubinstein, *Options Markets*, Prentice-Hall, Englewood Cliffs, NJ, 1985.
- [6] J., De Kleer, and J.S., Brown, "A Qualitative Physics Based on Confluences," *Artificial Intelligence*, 24:7–83, 1984.
- [7] J.E., Elton, and J.M., Gruber, *Modern Portfolio Theory and Investment Analysis (3rd edition)*, John Wiley & Sons, Inc., New York, NY, 1987.
- [8] D., Gentner, and A.S., Stevens, *Mental Models*, Erlbaum, Hillsdale, NJ, 1983.
- [9] J., Hull, *Options, Futures, & Other Derivative Securities*, Prentice-Hall, Englewood Cliffs, NJ, 1993.
- [10] B., Kuipers, and D., Berleant, "Using Incomplete Quantitative Knowledge in Qualitative Reasoning," *Proceedings of AAAI-88*, pp. 324-329, 1988.
- [11] B., Kuipers, "Qualitative Simulation," *Artificial Intelligence*, 29:289–338, 1986.
- [12] R., Saucedo, and E.E., Schiring, *Introduction to Continuous and Digital Control Systems*, The Macmillan Company, New York, NY, 1968.
- [13] B.C., Williams, "Interaction-Based Invention: Designing Novel Devices from First Principles," *Proceedings of AAAI-87*, pp. 349-356, 1987.

List of Figures

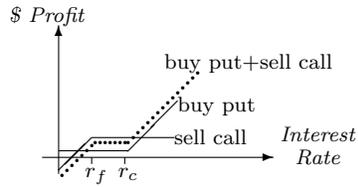
1. The configuration of a vehicle for the “cap-floor” payoff-profile
2. The configuration of a generic vehicle for the “cap” payoff-profile
3. The structure of a “buy put on bond” vehicle and its graphical representation
4. The transitions QSIM predicts for a “buy put option on bond” vehicle
5. Part of an ISA (specialization) hierarchy of instrument classes
6. A “ratio-spread” payoff-profile (and configuration)
7. Generic p-p’s abstracting the situation-specific p-p’s QSIM derives
8. Synthesizing a “butterfly-spread” using qualitative payoff-profiles
9. Two fundamental configurations of a system
10. QSYN’s synthesis of a ‘ratio-spread’ payoff-profile
11. Part of QSYN’s search tree for a ‘ratio-spread’
12. A quantitative analysis of one tradeoff between hedge vehicles

The goal loan rate contingent on the observed interest rate – avoid the risk that the loan rate will surpass l_c , and take the risk that the loan rate will decline below l_f .



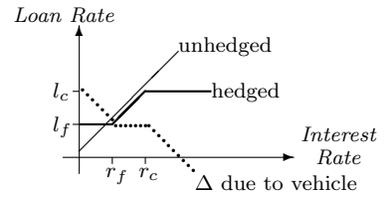
(a) A 'cap-floor' payoff-profile

The profit/loss pattern of the vehicle is the combined profit/loss patterns of 'buy one put' and 'sell one call' on bond B , where the put option has the higher strike price.



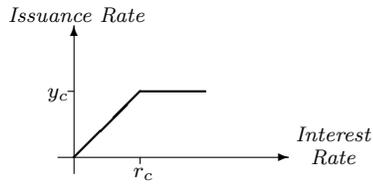
(b) A vehicle configuration

The vehicle's profit/loss pattern and the profit/loss pattern on the unhedged loan are combined to produce the goal payoff-profile.



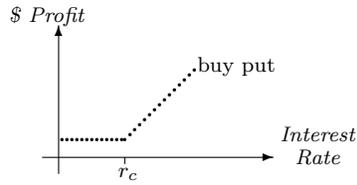
(c) A 'cap-floor' hedged position

The goal issuance rate contingent on the observed interest rate – avoid the risk that yield on issued bonds will surpass y_c , and preserve the ability to benefit from a lower interest rate.



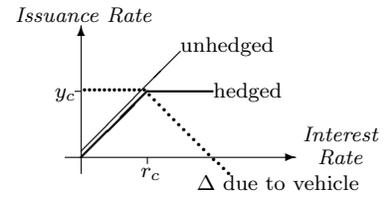
(a) A 'cap' payoff-profile

The profit/loss pattern of a 'buy one put on bonds' vehicle.



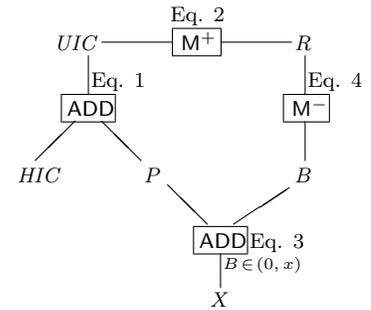
(b) A vehicle configuration

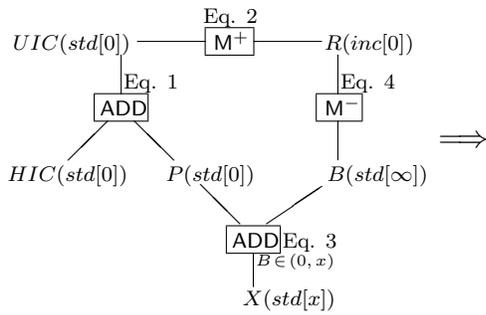
The vehicle's profit/loss pattern and the profit/loss pattern on the unhedged loan are combined to produce the goal payoff-profile.



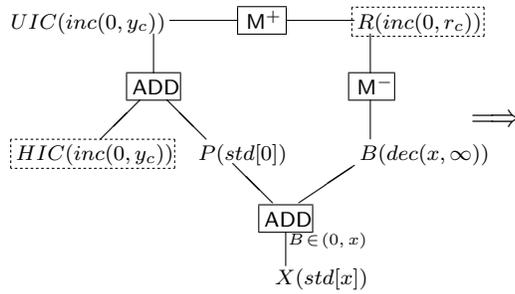
(c) A 'cap' hedged position

No	Qualitative Structural Equation	Explanation
1	$\text{ADD}(HIC, UIC, P)$	$HIC = UIC - P$. The hedged issuance cost (HIC) is the unhedged issuance cost (UIC) less the terminal value of the purchased put (P). This equation creates a 'buy' action affect.
2	$M^+(R, UIC)$	$R \propto^+ UIC$. The relationship between the risk-free interest rate (R) and the unhedged issuance cost (UIC).
3	$\text{ADD}(X, P, B)$ for $B \in (0, x)$	$P = \max(X - B, 0)$. The terminal value of a put (P) is zero when the bond price (B) is higher than the put's exercise price (x), or the difference between the exercise price and the bond price otherwise. Extracted from the valuation model for put options (i.e., Black-Sholes model [5]).
4	$M^-(R, B)$	$R \propto^- B$. The relationship between the risk-free interest rate (R) and the price of a bond (B), which is the put's underlying instrument. Extracted from the valuation model for bonds.

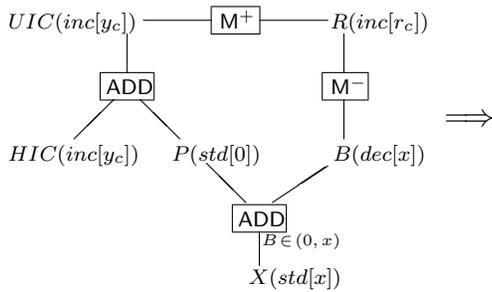




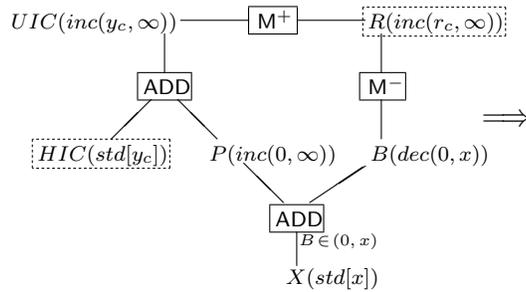
STATE 0: R is perturbed to start increasing



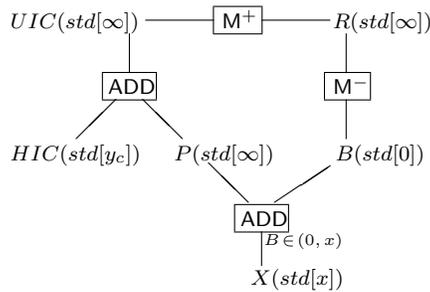
STATE 1: P remains zero as $Eq. 3$ is inactive



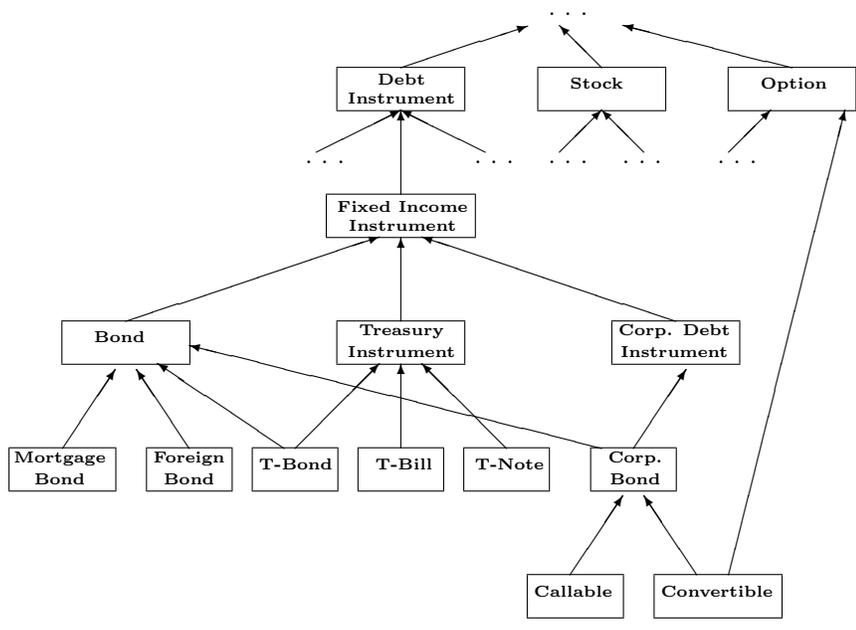
STATE 2: parameters keep on moving in their direction

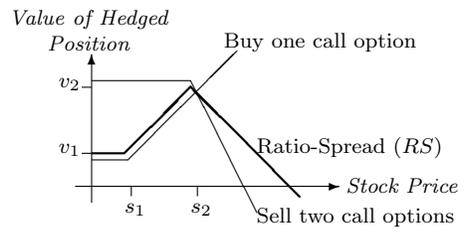


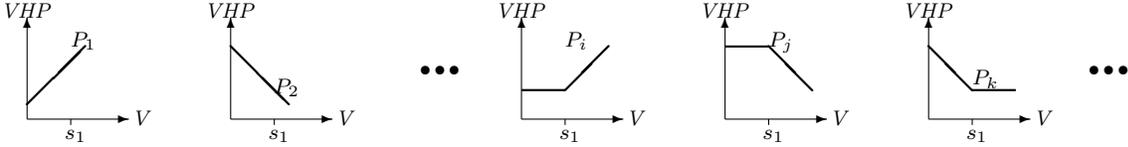
STATE 3: P starts increasing and offsets UIC 's increase



STATE 4: all parameters reach a 'steady' state

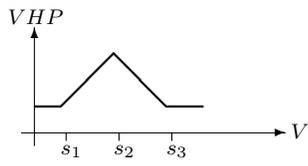




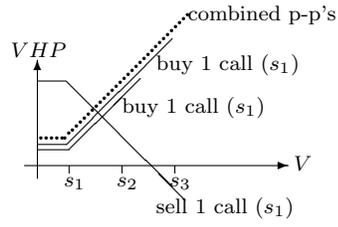


V - Hedged economic variable
 VHP - Value of hedged position
 s_1 - An arbitrary qualitative strike price

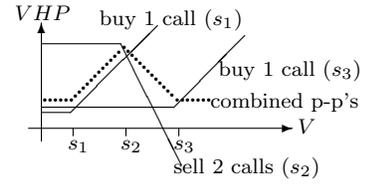
A 'butterfly' payoff-profile.



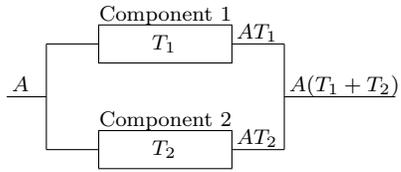
The combined qualitative p-p's do not produce a match even though they are supposed to.



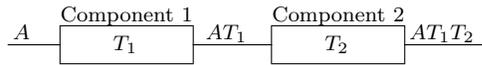
A match is produced only after the qualitative p-p's are stretched and/or steepened in a specific way.



T_i – transfer function of a component
 A – input to the system



(a) A system made from two components connected in parallel

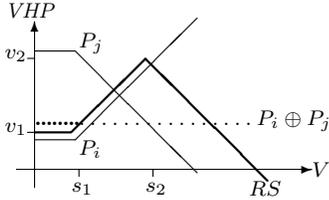


(b) A system made from two components connected in series

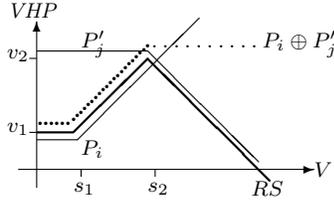
INPUT

$(P_i [(V \ 1 \ (0, s_1))(VHP \ 0 \ [v_1])][(V \ 1 \ (s_1, \infty))(VHP \ 1 \ (v_1, \infty))])$
 $(P_j [(V \ 1 \ (0, s_1))(VHP \ 0 \ [v_2])][(V \ 1 \ (s_1, \infty))(VHP \ -1 \ (-\infty, v_2))])$
 $(RS [(V \ 1 \ (0, s_1))(VHP \ 0 \ [v_1])][(V \ 1 \ (s_1, s_2))(VHP \ 1 \ (v_1, v_2))][(V \ 1 \ (s_2, \infty))(VHP \ -1 \ [-\infty, v_2))])$

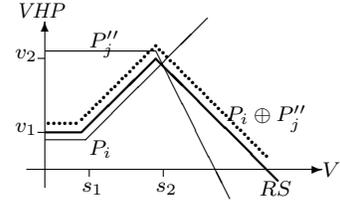
$$\begin{array}{l} (P_i[2] [(V * (s_1, \infty))(VHP \ 1 *)]) \\ (P_j[2] [(V * (s_1, \infty))(VHP \ -1 *)]) \\ \hline (\oplus [(V * (s_1, \infty))(VHP \ 0 *)]) \\ \text{No match with:} \\ (RS[2] [(V * (s_1, s_2))(VHP \ 1 *)]) \\ \hline \downarrow \\ \text{STRETCH}(P_j[1]) \\ \downarrow \\ (P_i[2] [(V * (s_1, \infty))(VHP \ 1 *)]) \\ (P'_j[1] [(V * (0, s_2))(VHP \ 0 *)]) \\ \hline (\oplus [(V * (s_1, s_2))(VHP \ 1 *)]) \\ \text{Match with:} \\ (RS[2] [(V * (s_1, s_2))(VHP \ 1 *)]) \end{array}$$

$$\begin{array}{l} (P_i[2] [(V * (s_1, \infty))(VHP \ 1 *)]) \\ (P'_j[2] [(V * (s_2, \infty))(VHP \ -1 *)]) \\ \hline (\oplus [(V * (s_2, \infty))(VHP \ 0 *)]) \\ \text{No match with:} \\ (RS[3] [(V * (s_2, \infty))(VHP \ -1 *)]) \\ \hline \downarrow \\ \text{STEEPEN}(P'_j[2]) \\ \downarrow \\ (P_i[2] [(V * (s_1, \infty))(VHP \ 1 *)]) \\ (P''_j[2] [(V * (s_2, \infty))(VHP \ -2 *)]) \\ \hline (\oplus [(V * (s_2, \infty))(VHP \ -1 *)]) \\ \text{Match with:} \\ (RS[3] [(V * (s_2, \infty))(VHP \ -1 *)]) \end{array}$$


(a) Start (match on 1st element of RS is found)



(b) Stretch 1st element of P_j to create P'_j (match on 2nd element of RS is found)



(c) Steepen 2nd element of P'_j to create P''_j (match on 3rd element of RS is found)

- △ P&L pattern of C_2 , if the price volatility of S increases
 - P&L pattern of C_1 , if the price volatility of S increases
 - P&L pattern of C_1 , if the price volatility of S does not change
 - P&L pattern of C_2 , if the price volatility of S does not change
 - P&L pattern of C_1 , if the price volatility of S declines
 - ◇ P&L pattern of C_2 , if the price volatility of S declines
- ⋯ qualitative P&L pattern of both options