INVESTMENT

Economics and finance view an agent as a system operating within one or more larger systems (e.g., markets). As in physical domains, this has allowed us to model economic and financial phenomena in terms of system composition, interrelatedness, components' interaction with the environment, and how components' behavior is controlled by decisions and policies that act as pumps, valves, and pressure regulators on the flow of funds in the pipes that connect components (1,2). Mathematical models developed based on this view aim to provide insight into the consequences of financial activities by translating decisions, events, and market forces into a language of cash flows.

Mathematical models have three general limitations. First, we must know the value of every parameter in a model before the model can be solved. In this sense, many financial models sometimes cannot be solved quantitatively because it is costly to acquire or develop precise estimates for their parameters (3). Second, intuitive reasoning with formal mathematics, as compared with prose, is difficult because of its limited interpretability (4). A mathematical model can neither explain its solutions nor the reasons for arriving at those solutions. For example, numeric simulation with such a model can predict a change in the behavior of a parameter, but it cannot explain what causes that change. Third, many mathematical models are inadequate for solving problems having combinatorially explosive search spaces. For example, models for optimizing investment portfolios involving many securities can be impractical even when used with parallel computers (5).

Because of these limitations, decision makers must rely on intuition and experience in reasoning about various financial phenomena, at least in early stages of the decision-making process (6). Interestingly, however, decision makers are often interested in merely understanding the qualitative nature of a problem before making decisions. For example, financial analysts usually translate large amounts of quantitative data into a few qualitative terms that are more insightful, which they can use to characterize a problem and subsequently select analytical techniques and/or generate solution alternatives (7,8). Thus, an early qualitative understanding of a problem is vital and largely determines, however implicitly, the alternatives considered. Yet, as research on human biases (9), human bounded rationality (10), and agency theory (11) indicates, decisions made based on intuition and experience are likely to be suboptimal.

In light of these observations, work on techniques of qualitative reasoning (QR)—an artificial intelligence (AI) approach to modeling and solving physics and engineering problems—aims to facilitate building knowledge-based systems (KBSs) that provide intelligent assistance to financial decision makers. QR techniques were originally developed to emulate humans' ability to reason intuitively about physical systems. A number of QR techniques have been used in several economic and financial KBSs, proving to be valuable in supporting various generic decision-making activities. These activities include

- Predicting economic behavior (12–14)
- Diagnosing deviations from a planned economic behavior (15,16)
- Explaining economic behavior (17–20)
- Planning actions to regulate economic and financial behavior (21,22)
- Configuring investment positions providing some goal behavior (13,20)

This article discusses the application of QR techniques in support of investment decision making. It first reviews some aspects of the investment process and its complexities. Then, it explains the way several QR techniques are used to overcome some of these complexities. Throughout the discussion, the article also points out the value of using QR techniques from an organizational and strategic perspective. The article concludes with a brief review of recent on-going research on QR techniques for financial and economics applications in general.

INVESTMENT AND QUALITATIVE REASONING

Investment is generally concerned with finding the best group of securities (i.e., position, portfolio) to hold, given properties of the available securities, the desired risk exposure and level of return, investor constraints and preferences, and the economic and legal environment (23). Figure 1 presents a top-down view of the investment process. The difficulty in this process is a result of the complexity inherent in two related subproblems: prediction and design.

Prediction entails identifying future economic trends and then selecting securities that stand to gain most from these trends (see Fig. 1). It starts with an assessment of the overall economy and its near-term outlook, to identity market trends (e.g., excess cash supply), risks one may seek to avoid, and movements in security prices. This assessment involves developing predictions about economic variables (e.g., money supply, interest rates) that directly affect the price, risk, and liquidity of securities or that just signal changes in future markets. These indicators, in turn, help to identify attractive market sectors (e.g., industries) or even specific firms whose securities are likely to have desired attributes (e.g., stocks with low price/earnings ratio). Relative to the securities issued by a specific firm, prediction can be viewed from two dual perspectives—"internal" and "external"—taken by the firm's management and by security analysts, respectively. Both perspectives study the financial actions of a firm (e.g., sales, borrowing). They try to relate the value of securities (e.g., stock, bond) the firm issued to the behavior of this firm and its economic environment. The "internal" perspective, in addition, focuses on understanding how the firm's past and present activities affect future strategic choices in the design

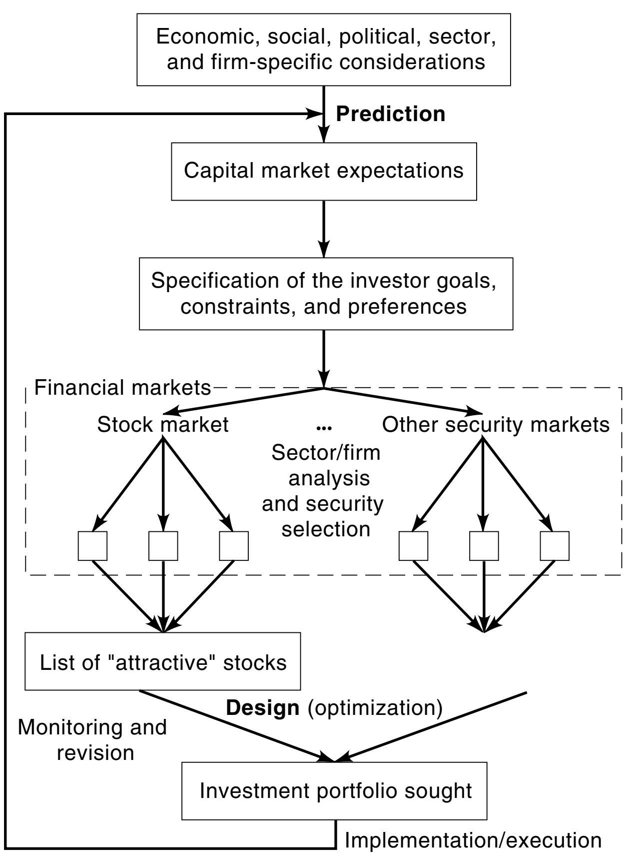

Figure 1. A top-down view of the investment process. The key complexities in this process are present in the prediction and design activities. *Prediction* is complex because of the growing universe of securities and their sophistication, as well as the uncertain, incomplete, and ambiguous data about this universe. *Design* is complex because of the combinatorial number of design alternatives associated with all security combinations and the possible proportions in which each security can be held.

of policies for regulating corporate behavior. Overall, developing and interpreting economic predictions is difficult for two reasons. One is the complexity brought about by the growing universe of securities (due to globalization) and the increasing sophistication of these securities. Another is the uncertain, incomplete, and ambiguous data available about this universe.

Design involves constructing a portfolio using the attractive securities identified (see Fig. 1). It usually entails multiobjective optimization. Constraints may exist for matching some goal risk profile (i.e., profit and loss pattern), matching the desired investment time horizon, not exceeding the available cash and credit in setting up the portfolio, and so on. The multiple objectives could be minimizing credit risk, minimizing setup cost, maximizing liquidity, and so on. Alternative positions are constructed and evaluated in light of the investor's objectives and constraints. Here, the major difficulty is due to combinatorics. The large universe of securities presents a tremendous choice in terms of how to design positions that exhibit some desired behavior. For investors, good predictions about economic factors cut down the risk associated with taking positions, but only to a limited degree. The real bottleneck here is complexity: when all security combinations and the possible proportions in which each security can be held are considered, the design problem is overwhelming.

Financial theories have yielded various formal models in support of both prediction and design. Many of these models have been extensively tested, and so investment specialists find them appealing because of their credibility. Yet, surveys show that, because these models have certain limitations, they are not used extensively, especially in the early stages of the prediction and design activities (24). For example, because the combinatorics involved in designing portfolios is prohibitive, most portfolio optimization models are impractical even when used with parallel computers. Consequently, investment specialists are often forced to rely extensively on heuristics embodying insights and perceptions that they have gained over years of experience. Unfortunately, as we indicated earlier, decisions made largely based on experiential heuristics are likely to be suboptimal.

Research on QR techniques focuses on enabling the development of KBSs that can help to leverage formal models, for example, by faciliating their use with incomplete, inconsistent, and imprecise data. In what follows, we present three general examples where QR techniques are used to deal with the investment complexities involved in prediction and design. Each of these examples helps to see more specific complexities, the significance of these complexities in light of limitations of formal models, and the way that these limitations of formal models are avoided using different QR techniques.

ANALYZING CORPORATE BEHAVIOR

Financial managers are usually interested in controlling effects of the economic environment on corporate behavior and in turn on securities issued by their firm. Doing so first requires understanding what causes corporate behavior and how it comes about, and then making strategic choices in the design of corporate policies that "improve" this behavior (1).

Limitations of Corporate Planning Models

Controlling corporate behavior by making strategic choices requires understanding how corporate behavior results from corporate structure. The complexity faced here is a result of the many parameters and relationships characterizing a firm. It is difficult to trace how these parameters interact to produce the overall corporate behavior. Unfortunately, formal financial models simply cannot capture the volume of relationships between these parameters. Even the most powerful models are suitable for analyzing only single pieces of the puzzle. As Brealey and Myers (3, p. 683) explain, “There is no model or procedure that comprehends all the complexity and intangibles encountered.”

Decision makers must therefore rely on intuition and experience in assessing the consequences of strategic choices and policies. However, because the human mind is simply incapable of evaluating the implications of more than just a few interactions between parameters (1), understanding how corporate behavior results from its structure without the aid of automated tools can be time consuming and erroneous.

Formal models aimed at helping to handle this complexity focus on providing simulation of the enterprise (1). They allow financial managers to probe the solution space of a problem so as to gain insight beyond the mere solution of a model, until a level of understanding is reached that would support making a decision. Simulation typically involves an iterative process: perturb model, identify impacts on performance measures, and design policies to regulate behavior. This process involves what-if analysis that adaptively explores a problem by performing a preconceived set of runs that test the effect of various strategic choices. Simulation is most effective when the modeler understands why a particular structure produced the simulated behavior. Unfortunately, conventional simulation cannot explain its solutions nor the way it arrives at these solutions. The interpretation of, and insight drawn from, generated data are left to the decision maker. In effect, simulation does not even tell which alternatives are worth examining. These limitations are compounded by the fact that quantitative simulation cannot be directly used for problems involving parameters whose precise value is unknown.

Interestingly, however, decision makers are often concerned with merely understanding the qualitative characteristics of a problem, especially in the early stages of the decision-making process. In some cases a qualitative understanding is sufficient to make a decision, whereas in other cases it is simply a prerequisite to the design and/or selection of suitable formal models and their solution using mathematical techniques (7,8). In either case, it largely determines, however implicitly, the alternative strategic choices considered.

It has been shown that QR techniques can help decision makers develop such an understanding. For example, a QR technique called *qualitative simulation* is capable of reasoning with imprecise knowledge and thus can help to develop qualitative insights into a complex problem in the early stages of the decision-making process. The motivation behind using QR techniques to analyze corporate behavior is grounded in the realization that a firm is conceptually viewed as a system. This view has allowed us to model financial phenomena mathematically in terms of system composition, interrelatedness, and components' interaction with their environment.

Qualitative Simulation

Qualitative simulation (QSIM) is the most general and commonly used QR technique (25,26). As in other QR techniques, QSIM's approach is anchored in the recognition that humans use a qualitative causal calculus to reason about the behavior of physical systems (27). QSIM can derive the qualitative behavior of a system based on that system's structure as well as explain this behavior in intuitive terms. The key ideas behind how QSIM works follow:

1. The structure of a system is described by structural equations modeling connections between its characterizing parameters.
2. Given that the system is in some initial state, a change in the state of parameters propagates locally to other parameters through structural connections.
3. The qualitative behavior of a parameter is described by the transitions it makes from one state to another.
4. The qualitative behavior of a system is described by the interaction of behaviors of its characterizing parameters.

Structure is described in terms of components and their connections. A component is modeled by one or more real-valued parameters (continuous functions), each associated with a finite set of *landmark values*—points where something special happens to the parameter (e.g., an extremum). A structural connection is modeled by a qualitative constraint equation that restrict the values that the parameters can take on. QSIM reasons with two types of qualitative constraint equations. One type is for specifying simple mathematical relationships [i.e., addition (ADD), multiplication (MULT), derivative of time (DERIV), and unary negative (MINUS)]. Another type is for specifying functional relationships between parameters [i.e., a monotonic change of two parameters in the same direction (M^+) or in opposite directions (M^-)]. These constraints are useful when the precise value of constants that relate parameters is difficult or costly to measure [e.g., $Y = kX$, where k is a constant, is represented as ($M^+ YX$)]. Part of the description of structure includes information about the correspondence of landmark values across connected parameters.

Given the structure of a system and assuming that the system is perturbed, QSIM generates all the behaviors of that system and represents them using a *transition graph*. In this graph, each node represents the qualitative state of the system at a specific time point, every pair of adjacent nodes represent two temporally adjacent qualitative states, and every path from the initial state node through the graph represents one behavior of the system. Each qualitatively distinct state of the system (represented by a node) is described by the qualitative state of every system parameter at one specific *distinguished time point*, a point where something special happens to the system. A *qualitative state* of a parameter is the pair $\langle qdir, qval \rangle$, where $qdir \in$ {decreasing = − 1, steady = 0, decreasing = 1} is the direction of change of the parameter value over a *qval*—a point corresponding to a landmark value or an interval between two landmark values.

When one or more of the parameters of a system in equilibrium are perturbed, QSIM propagates the change to other parameters so as to derive the next qualitative state of every parameter and of the system as a whole. QSIM continues to propagate change in this fashion, until all parameters reach a steady state or a boundary *qval* or exhibit a cyclic behavior. Specifically, this simulation process involves the following iterative steps (26):

- *Identify for each parameter its potential transitions.* Because each parameter is a continuously differentiable function, theorems from calculus restrict the moves that the function can make from one point to another. For instance, if the derivative of a function is positive over (x_i, x_{i+1}) $\in \Re$, it must become zero at x_{i+1} before it can become negative. Thus, the next potential transitions of a parameter are selected from a finite set of legal transitions it can have from any one state to another. To illustrate, if the current state of parameter X is $\langle$inc $(x_1, x_2)\rangle$, the potential transitions for X are $\{\langle$std $[x_2]\rangle$, $\langle$inc $[x_2]\rangle$, $\langle$inc $(x_1, x_2)\rangle$, $\langle$std $[x^*]\rangle\}$. The last transition represents a case where the previously unknown landmark value x^* ($x_1 < x^* < x_2$) is discovered by QSIM (as a result of having constraints that force X to become steady).
- *Filter the potential transitions for each parameter.* Eliminate combinations of transitions that are inconsistent with the system's structure. For instance, the constraint ADD(X, Y, Z) does not allow for both X and Y to be increasing while Z is steady. This filtering process finds only the possible transitions. For each consistent set of parameters' qualitative stats found, QSIM adds a node to the transition graph to represent the next qualitative state of the system.
- *Characterizes each derived next state of the system.* The next state can be either (1) an equilibrium (quiescent) state where all parameters are steady, (2) a state indicating a cyclic behavior (a state identical to some previous state), (3) a state indicating a divergent behavior (i.e., one or more parameters go to $\pm\infty$), or (4) a state indicating that one or more parameters are still changing (i.e., moving toward a landmark value).

Step 2 alludes to a key issue related to the pruning of "unreal" behaviors. Each path in the transition graph represents one possible behavior. In some cases, however, a path may represent a spurious behavior. Because parameters are characterized only qualitatively, sometimes there is insufficient information to determine the behavior of parameters that are affected by competing tendencies. For example, consider the constraint $X + Y = Z$. If at some time point X is increasing and Y is decreasing over the same interval in $\Re$, the behavior of Z is ambiguous. Therefore, QSIM creates a branch in the graph to account for the three possible behaviors of Z—steady, increasing, or decreasing (for $X = Y$, $X > Y$, and $X < Y$, respectively). The possibility that this ambiguity may never arise in reality implies that two of these alternatives lead to spurious behaviors. A variety of methods have been developed to help prune spurious behaviors (27). For example, one requires the use of knowledge about the sign of higher derivatives, whereas another incorporates numeric information whenever an ambiguity arises.

Predicting Qualitative Consequences of Policies

The modeling of problems for use with QR techniques is anchored in a systemic view of corporate structure. This structure is modeled in terms of accounting relationships between the various parameters characterizing an enterprise. Changes in the behavior of parameters are modeled by changes in flow accumulation of funds in various fund sinks, where this behavior is regulated by decisions and policies that act as pumps, valves, and pressure regulators on flow of funds in pipes connecting parameters.

The following simplified problem illustrates how the use of QSIM can help in making strategic choices. [Applications involving larger problems are discussed in Ref. (12).]

Scenario 1: Trust Ltd. is a publicly traded firm that uses one part of its net operating income (NOI), the retained earnings (RE), to finance a new project and the other part, the allocated dividends (AD), to pay dividends to its shareholders. Unless NOI and RE change, the AD = ad and the amount of dividend per share (DPS = d) remain constant over time. The value of one part of the firm's assets, the equity E, equals the number of common stock shares (CS = cs) multiplied by the stock price (P = p). The other part of the firm's assets is debt (D).

The firm is considering ways to "improve" its image as a high-profit firm. One idea is to temporarily increase the amount paid as DPS, without increasing D. Starting at some time point t_0, DPS is to increase from its current level d to a new level d^*, for a short period of time ending at t_1. At t_1, DPS is to be reduced back to a level that maintains the amount of AD prior to the increase in DPS. The goal is to predict the effects of this policy as well as understand what causes these effects.

Based on this description, the qualitative constraint equations in Fig. 2(a) describe the structure of the system (firm) under examination. Figure 2(b) offers a graphical representation of the system's structure, to help the reader trace QSIM's simulation results. Given that the system is initially in equilibrium, its DPS is perturbed in a specific fashion, and the goal is to understand how and why the system reacts to the change in DPS. This is the goal of Trust's management who seeks to understand behavior qualitatively, before considering the level of increase in DPS and its duration.

Solution 1: QSIM produces a transition graph. In this graph, one path ends with a state where, after a long increase in DPS, parameters reach a state of divergence—stock price P goes to 0 and the number of common stocks goes to ∞. Another path, whose nodes are described in the table in Fig. 2(c), ends with a state where, after a short increase in DPS, the system reaches a state of equilibrium—DPS drops below its initial level d and then stabilizes at a lower level d_*, the number of common stocks stabilizes at cs^* above the initial number of stocks cs, and the stock price P stabilizes at p_* below the initial stock price p. QSIM explains the behavior that leads to equilibrium as follows:

- During (t_0, t_1), as DPS increases from d to d^*, the global dividends paid (GDP) exceeds AD, causing CD to become positive (a cash deficit), which in turn causes the number of common stocks issued (CSI) to become positive (issue common stocks) and CS to increase.
- As E is steady and CS is increasing, P starts to decline below p. At t_1, DPS reaches d^* and starts declining and reaches d at t_2. At t_2, as DPS declines below d, GDP reaches a pick level (gdp^*) and starts declining toward gdp (= ad), causing CSI to pick at t_3, and to start declining.

```
(defnet DIVIDEND-POLICY
  (functions     (DPS GDP CD CSI CS P F E D RE NOI AD))
  (constraints   (add  E D F)      % Firm value = Equity + Debt
                 (mult CS P E)        % E = Common-Stocks * stock-Price
                 (add RE AD NOI)      % Net-Operating-Income = Retained-Earnings + Allocated-Dividends
                 (add  AD CD GDP)     % Allocated-Dividends = Cash-Deficit + Global-Dividend-Paid
                 (mult CD CSI P)      % Cash-Deficit = Common-Stocks-Issued * stock-Price
                 (mult CS D GDP)      % Global-Dividend-Paid = Common-Stocks * Dividend-Per-Share
                 (d/dt CS CSI))       % dCS/dt = Δ* CS = CSI
  (landmarks     (DPS  (minfinity 0 d d* infinity))
                 (GDP  (minfinity 0 gdp infinity))
                 (CD   (minfinity 0 infinity))
                 (CSI  (minfinity 0 infinity))
                 (CS   (minfinity 0 cs infinity)
                 (P    (minfinity 0 p infinity)))
  (ranges        (F    (f *constant*))
                 (E    (e *constant*))
                 (D (d *constant*))
                 (NOI  (noi *constant*))
                 (AD   (da *constant*))
                 (P    ((0 inf) nil)))
  (initialize    (DPS  (inc (d d*))))))
```

(a)

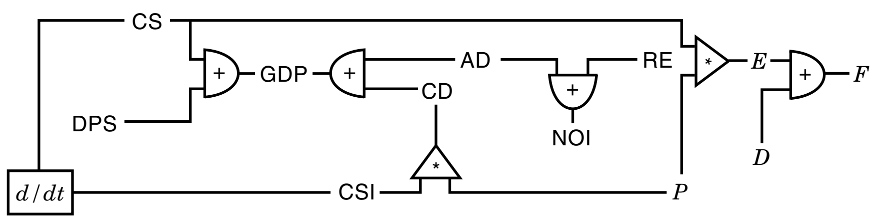

(b)

Time	DPS	GDP	CD	CSI	CS	P	Explanation
(t_0, t_1)	$\langle$inc $(d, d^*)\rangle$	$\langle$inc $(gdp, \infty)\rangle$	$\langle$inc $(0, \infty)\rangle$	$\langle$inc $(0, \infty)\rangle$	$\langle$inc $(cs, \infty)\rangle$	$\langle$dec $(0, p)\rangle$	policy
$[t_1]$	$\langle$std $(d, d^*)\rangle$	$\langle$inc $(gdp, \infty)\rangle$	$\langle$inc $(0, \infty)\rangle$	$\langle$inc $(0, \infty)\rangle$	$\langle$inc $(cs, \infty)\rangle$	$\langle$dec $(0, p)\rangle$	policy
(t_1, t_2)	$\langle$dec $(d, d^*)\rangle$	$\langle$inc $(gdp, \infty)\rangle$	$\langle$inc $(0, \infty)\rangle$	$\langle$inc $(0, \infty)\rangle$	$\langle$inc $(cs, \infty)\rangle$	$\langle$dec $(0, p)\rangle$	policy
$[t_2]$	$\langle$dec $d\rangle$	$\langle$std $gdp^*\rangle$	$\langle$std $cd^*\rangle$	$\langle$inc $(0, \infty)\rangle$	$\langle$inc $(cs, \infty)\rangle$	$\langle$dec $(0, p)\rangle$	change state
(t_2, t_3)	$\langle$dec $d\rangle$	$\langle$dec $(ad, gdp^*)\rangle$	$\langle$dec $(0, cd^*)\rangle$	$\langle$inc $(0, \infty)\rangle$	$\langle$inc $(cs, \infty)\rangle$	$\langle$dec $(0, p)\rangle$	change state
$[t_3]$	$\langle$dec $d\rangle$	$\langle$dec $(ad, gdp^*)\rangle$	$\langle$dec $(0, cd^*)\rangle$	$\langle$std $csi^*\rangle$	$\langle$inc $(cs, \infty)\rangle$	$\langle$dec $(0, p)\rangle$	change state
(t_3, t_4)	$\langle$dec $d\rangle$	$\langle$dec $(ad, gdp^*)\rangle$	$\langle$dec $(0, cd^*)\rangle$	$\langle$dec $(0, csi^*)\rangle$	$\langle$inc $(cs, \infty)\rangle$	$\langle$dec $(0, p)\rangle$	change state
$[t_4]$	$\langle$std $d_*\rangle$	$\langle$std $ad\rangle$	$\langle$dec 0$\rangle$	$\langle$std 0$\rangle$	$\langle$std $cs^*\rangle$	$\langle$std $p_*\rangle$	equilibrium state

(c)

Figure 2. QSIM's application to the dividend policy problem. (a) A specification of the problem using QSIM's constraint equations formalism. (b) A graphical representation of the constraint equations can help the reader to manually trace QSIM's simulation. (c) The qualitative states on one of the paths in QSIM's transition graph indicates that the system reaches a new equilibrium—DPS drops and then stabilizes at d^* below the initial dividend per share level d, CS stabilizes at cs^* above the initial number of common stocks cs, and P stabilizes at p^* below the initial stock price p.

- At t_4, when GDP becomes steady, CD and CSI reach zero and become steady, causing CS to reach a pick level cs^* and to become steady, which in turn causes P to reach a lower level p_* and become steady. Because all parameters become steady, QSIM concludes that the system reached a new equilibrium.

This same scenario can be analyzed from the "external" perspective, by an independent security analyst who wants to identify how certain publicly announced corporate policies affect the value of securities issued by the corporation. The analyst would use QSIM to conduct the same analysis summarized previously. Alternately, we can think of intelligent programs that intercept a live news wire to read and interpret news in order to detect qualitative changes in economic variables like trade balances and government expenditure (28) and in turn activate QSIM on prestored models in order to identify interesting market events (12). These programs would act as "bell ringers" that can have a strategic impact on the ability of a financial institution to rapidly react to the news signaling market changes.

MAKING ACTUAL INVESTMENT DECISIONS

We saw how QR techniques can help in financial analysis, for example, for the purpose of assessing the value and sensitivity of securities to corporate policies and economic changes. We next show how these techniques can also assist in making actual investment decisions based on such qualitative assessments.

Complexities in the Design of Positions

Investment specialists are usually interested in designing portfolios (i.e., combinations of stocks, options, bonds, or future contracts) that exploit profit opportunities in the market place and meet the investor's requirements. In principle, formal financial models aim to support this design endeavor, for example, by helping to understand how the behavior of a portfolio results from its structure. Because investment involves complex strategies, where returns on most strategies are contingent on future uncertain market states, such an understanding is vital to the ability to design portfolios that are robust to deviations from forecasted economic trends (23).

As with corporate planning models, formal design models have two key limitations. First, they cannot explain their results. Investment specialists typically use quantitative what-if analysis to understand the contingent nature of returns and their effect on the value (behavior) of a portfolio, in terms of the value of its components and their relationships to economic parameters. More importantly, these models often cannot handle the large number of investment possibilities and their sophistication. McInnes and Carlton (6, p. 568) explain: "Computationally, an exhaustive analysis of all the possible investment combinations rapidly becomes intractable as the number of investment programs increases. Human judgment has to intervene to reduce the number of possibilities to be explored by formal analysis to a manageable set." Yet, as we mentioned earlier, because of human cognitive limitations, unaided analysis in the early design stages can have critical implications on later stages. This problem is magnified by the fact that investment specialists typically specialize only in subsets of the many types of securities that can be used to construct portfolios. This exposes them to a tunnel vision problem that leads to suboptimal investment decisions.

The last problem is aggravated by the current tendency of financial institutions to gain a strategic advantage by moving toward integration, as more information is becoming available about securities traded in domestic and global markets. Under this scenario, investment specialists would seek to design portfolios that exploit intricate opportunities present in the marketplace, as long as there are intelligent tools to help them manage the additional complexity brought about by considering a larger set of securities. For the most part, such tools need to do a lot of screening and to present only the most promising alternatives for further quantitative analysis. Of course, such tools must first be able to configure automatically alternative portfolios that meet certain investor requirements which are usually specified qualitatively.

In light of these complexities and limitations of quantitative design models, it seems that QR techniques can play an important support role in the design endeavor. Two factors indicate that these techniques can be used for this purpose. First, qualitative abstraction is a powerful means that investment specialists use to cope with the complexity involved in assessing the large number of investment possibilities (7,8). Additionally, a systemic view of the design endeavor can be used here as well, because investment models are usually developed based on principles from cybernetics and control theory [e.g., the Black-Scholse model (23)].

Designing Simple Positions

The next small example illustrates how QR techniques can be useful in the early design stages, where alternative portfolios are configured, prior to their extensive evaluation using quantitative analysis.

Scenario 2: Trust Ltd. decided to finance its new project using a floating rate long-term loan, tied to the 6 month Eurodollar rate. The loan rate is 7%—current $4\frac{1}{2}$% LIBOR (London InterBank Offer Rates) plus $2\frac{1}{2}$% stamping fee. Trust's management believes that there is a good chance that the risk-free interest rate would rise in the next 6 months, and this can significantly affect the cost of the loan. At the same time, Trust's management believes that there is a possibility that the risk-free interest rate will decline, in which case the interest rate paid on the loan will decline as well. Trust's management hence seeks to protect against the risk of increase, while preserving the ability to benefit from a decline, in interest rates.

In this scenario, Trust Ltd. seeks to *hedge* interest rate risk, by holding a position that has the "cap" risk profile seen in Fig. 3(d). *Hedging* is an investment problem that is concerned with the design of controls that minimize the adverse affects of possible losses or their consequences (29). In this article, we consider hedging to deal only with controls that involve the purchase and/or sale of securities, not actions concerning real assets (e.g., relocating production facilities to the foreign markets where finished goods are sold in order to avoid foreign exchange risk).

Investment positions having the cap risk profile can be configured using cash securities and their derivatives. These securities include Treasury securities (T-bond, T-notes, and T-bills); futures on Eurodollar securities (i.e., dollar deposits outside the United States); future contracts on LIBOR; and call and put options on the previously mentioned securities, on short-term and long-term interest rate, and on the MUNI (Municipal) bond index. One specific position is explained here.

Solution 2: Trust Ltd. can purchase put options on some bond B with strike price b. A *put option* on B provides its buyer the right to sell, and obligation its seller to buy, units of B for an agreed-upon strike price b at some future expiration date. An increase in interest rate will cause the price of B to decline below b, allowing the firm to profit from selling bonds for b and to thus offset the extra cost paid for the loan. Alternately, a decline in interest rate will make the put option worthless but allow the firm to benefit from a lower loan cost that offsets, and more, the purchase cost of the put option.

Parameters	
LC—loan cost	R—interest rate
HLC—hedged loan cost	B—bond value
P—value of put option	X—strike price of the put option (constant)

Constraints	Explanation
1. ADD(HLC, LC, P)	HLC = LC − P
2. $M^+(R, \text{LC})$	$R \propto^+ \text{LC}$
3. ADD(X, P, $B \in (0, x)$)	$P = \max(X - B, 0)$
4. $M^-(R, B)$	$R \propto^- B$

(a)

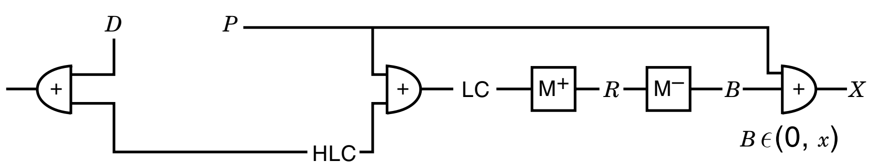

(b)

	R	LC	B	X	P	HLC
Initial state	⟨inc [0]⟩	⟨std [ε]⟩	⟨std [∞]⟩	⟨std [x]⟩	⟨std [0]⟩	⟨std [ε]⟩
State 1	**⟨inc [0, r_c]⟩**	⟨inc (ε, l_c)⟩	⟨dec (x, ∞)⟩	⟨std [x]⟩	⟨std [0]⟩	**⟨inc (ε, l_c)⟩**
State 2	⟨inc [r_c]⟩	⟨inc l_c)⟩	⟨dec [x]⟩	⟨std [x]⟩	⟨std [0]⟩	⟨inc [l_c]⟩
State 3	**⟨inc (r_c, ∞)⟩**	⟨inc (l_c, ∞)⟩	⟨dec (0, x)⟩	⟨std [x]⟩	⟨inc (0, ∞)⟩	**⟨std [l_c]⟩**
Terminal state	⟨std [∞]⟩	⟨std [∞]⟩	⟨std [0]⟩	⟨std [x]⟩	⟨std [∞]⟩	⟨std [l_c]⟩

(c)

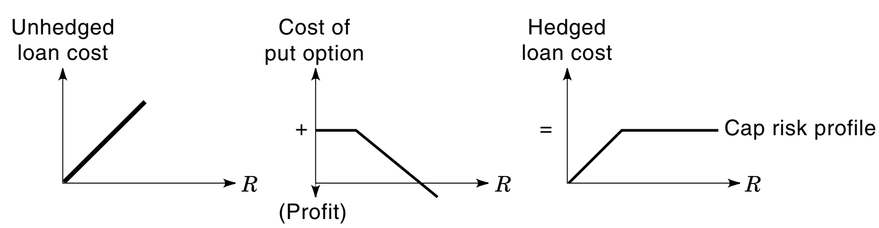

(d)

Figure 3. QSIM configures a position with a cap risk profile. (a) The parameters and qualitative constraint equations provided as input to QSIM. (b) A graphical representation of the constraint equations can help the reader to manually trace QSIM's simulation. (c) Qualitative states on one of the paths in QSIM's transition graph, where states 1 and 3 constitute the derived risk profile. (d) Graphically plotting the derived risk profile shows that it matches the goal (cap) risk profile.

How can QR techniques help configure such a position? When the composition of a position is known, QSIM can derive the position's risk profile (behavior) under the market scenario of concern and compare it against the goal risk profile. The simplest position is one containing a single component (security) in addition to the asset being hedged (e.g., loan). Its "structure" is described by two things. One is the structural equation POS = UA ± S, stating that the value of the POSition is the value of the Unhedged Asset plus (minus) the value of the security sold (purchased). The other thing is the *valuation model* of the security purchased or sold. Causal relationships between economic variables and the value of a specific security are each modeled by a structural equation that specifies how a certain economic variable affects the value of that security (23). The set of equations modeling these relationships for a specific security is called a valuation model. This model's analytic solution, called the *pricing model*, is used to compute the fair market price of that security. Because different types of securities are sensitive to different sets of economic variables, they each have a different valuation model.

To illustrate how QSIM derives the scenario-specific risk profile of a position, consider the example of using a "purchase put option on bond" position to cap Trust's loan cost. The cap risk profile is expressed symbolically as the sequence of pairs:

$$\{[(R\langle\text{inc}(0, r_c)\rangle)(\text{HLC}\langle\text{inc}(0, l_c)\rangle)]\ [(R\langle\text{inc}(r_c, \infty)\rangle)(\text{HLC}\langle\text{std}[l_c]\rangle)]\}$$

where R is the risk-free interest rate, HLC is the hedged loan cost, and r_c is the risk-free interest rate level corresponding to the cap level l_c on the cost. The input for QSIM includes the qualitative structural equations in Fig. 3(a), and the initial state of every parameter when R is zero. To allow the reader to trace QSIM's simulation results with greater ease, Fig. 3(b) offers a graphical representation of the qualitative structural equations describing the system's structure. In the initial state, interest rate is zero, the price of a yield-bearing bond is high and positive (infinite in the limit), the value of a put on that bond is zero, and the loan cost (hedged and unhedged) is an infinitesimally small ϵ (because theoretically a firm can offer to pay little interest to get the loan). We then run QSIM upon letting R increase over the range $(0, \infty)$. A trace of the states QSIM derives is presented in Figure 3(c) (ignoring the time dimension for simplicity). In state 1, R's increase causes B to start declining and LC to start increasing, in compliance with constraints 4 and 2, respectively. Because B has not yet reached x, the put's strike price P remains zero, complying with constraint 3, and HLC starts increasing to comply with LCs increase in constraint 1. In state 3, as R continues to rise, B declines below x, and P begins to increase in compliance with constraint 3. In turn, HLC becomes steady at the cap level l_c because QSIM is pre-told to assume that the increase in P balances off LC's increase in constraint 1. This assumption is based on the notion that a hedging position is constructed to balance off changes in the value of the unhedged position. This is possible by controlling in later design stages the precise number of units of the security purchased/sold. The risk profile of the position being analyzed is embedded in the sequence of states QSIM derives. These states are printed in bold in the table in Fig. 3(c). A comparison of this derived risk profile with the goal risk profile would thus conclude that a "purchase put option on bond" position can cap Trusts' loan cost [see Fig. 3(d)].

Pragmatic Issues

Configuring all one-security positions that provide the goal risk profile requires applying QSIM for every individual security available in the marketplace in the fashion described previously. The computational intensity this involves can be inhibitive. To deal with this problem, we can rely on other QR techniques and exploit domain-specific heuristics.

Qualitative abstraction over knowledge about securities can limit the application of QSIM needlessly. Specifically, securities naturally fall into classes, forming a specialization hierarchy like the one that follows:

```
(Security
    (Debt-Security
        (Fixed-Income-Security
            (Treasury-Security (T-Bill  T-Bond
                T-Note  ...))
            (Bond              (Mortgage-Bond
                T-Bond  Foreign-Bond  ...))
            (...))
        (Corporate-Debt-Security
            (Corporate-Bond    (Callable-Bond
                Convertible-Bond  ...))
            (...))
        (...))
     (Stock            (...))
     (Option           (...))
     (Future-Contract  (...))
     (...))
```

Because all securities of the same class have the same valuation model, the natural grouping of securities can be used in two ways. First, QSIM can be applied collectively for all securities of the same class. Second, QSIM can be applied only for each class of securities whose valuation model is a generalization of the valuation models of other security classes. In the specialization hierarchy, the qualitative valuation model of one class of securities can be a specialization of the valuation model of other classes. For example, the valuation model of bond options is a specialization of the Black-Scholes model used to derive the valuation model of other types of options (30).

Two sample heuristics for making the use of QSIM even more tractable follow. First, because the sale/purchase of a security that is insensitive to the relevant risky variable (e.g., interest rate) is meaningless from the standpoint of Trusts Ltd., QSIM must be applied only for security classes whose valuation model references this variable. Second, as the risk profiles derived upon buying and selling a certain security are symmetrical (because investment is a zero-sum game), QSIM can be used to derive the risk profile only for selling a security.

DESIGNING COMPLEX INVESTMENT POSITIONS

Realistically, the goal risk profile of an investor such as Trust Ltd. can be more complex, in which case it is necessary to configure multisecurity positions. For example, consider the long-term loan we discussed in scenario 2. Suppose that after rethinking the opportunities that a hedge position can provide, Trust's management agrees that the interest rate is not likely to drop below 6%. Like in scenario 2, Trust wants to "cap" the loan cost at a level that corresponds to an 8% interest rate. In addition, Trust seeks to set a "floor" on the loan cost at a level that corresponds to a 6% interest rate, by selling securities to another investor who believes that the interest rate will drop below 6%. Should the interest rate remain above 6%, Trust's profit would be what it receives for the securities it sells; otherwise, Trust's loss would be what it could save from paying less than 6% interest rate on its loan. This rather speculative investment behavior that Trust's management is exhibiting might seem unusual. However, by now most sophisticated firms are using the notion of hedging not just to protect against loss but also to generate profits based on their understanding of the marketplace.

The next scenario illustrates the role of QR techniques in configuring multisecurity positions. This scenario parallels the one Trust's management is facing, although it involves stock options. It is easier to understand how to configure multisecurity positions for hedging fluctuations in stock prices, instead of fluctuations in the cost of a loan that are brought about by fluctuations in interest rate. In other words, where the function $H(\ )$ denotes the value of a hedge position, it is easier to look at a case involving $H(stock)$ rather than $H(loan(interest\text{-}rate))$.

Scenario 3: An investor speculates that over the next six months the price of some stock S will rise above s_L but not above s_M. The investor decides that if S rises to somewhere between s_L and s_M, he would like to make a profit; and, if S rises above s_M, he is willing to take a loss with a limit that

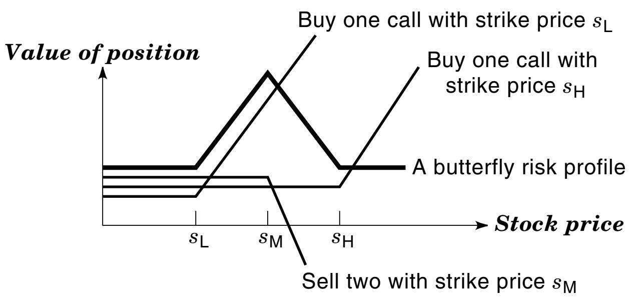

Figure 4. A butterfly risk profile. Configuring a multisecurity position with a butterfly-like risk profile involves searching the explosive space of linear combinations of elementary risk profiles. Abstracting all elementary risk profiles similar in shape into one qualitative risk profile drastically lowers the number of risk profile permutations, but it also results in the loss of useful information. We cannot synthesize a butterfly using two similar elementary risk profiles, corresponding to the purchase of call options with strike prices s_L and s_H, unless we can rediscover lost information about the ordering of strike prices.

corresponds to price s_H. Accordingly, he defines a goal risk profile called "butterfly" (see Fig. 4).

Solution 3: One way to derive this risk profile is to trade *call options* on the stock—buy a call option with strike price s_L, sell two calls with strike price s_M, and buy a call with strike price s_H. A *call option* provides its buyer with the right to purchase, and its seller with obligation to sell, shares of the stock for an agreed upon strike price at a future expiration date. Based on this definition, the options' combination works as follows:

- If $s_L \leq S \leq s_M$, the investor will profit from the cash received for the two calls sold and from exercising the purchased call with strike price s_L (i.e., buying shares for s_L).
- If $s_M \leq S \leq s_H$, the investor will gain on the purchased call with strike price s_L and lose on the purchased call with strike price s_H and on the calls sold (i.e., selling shares for s_M).
- If $S > s_H$, the investor's gain on the calls purchased will offset the loss on the calls sold.

Role of QR Techniques

This example shows that a position is described in terms of the securities purchased/sold and their unit proportions. This description is derived from how the risk profile of a position is composed as the algebraic sum of elementary risk profiles that are each associated with buying/selling one or more units of a specific security.

Given the elementary risk profiles of individual securities, configuring positions involves searching for all linear combinations of elementary risk profiles that match a goal risk profile. In other words, given a set of elementary piecewise linear functions with edges having a real-valued slope over some range in $(-\infty, \infty)$, the problem is to find all ways to synthesize a goal function using these elementary functions. This search problem is subject to combinatorial explosion. Considering only option-based securities, the number of possible permutations of risk profiles is 2^{4n}, where 4 stands for the risk profiles of sell call, buy call, sell put, and buy put, and n is the number of different strike prices (n can be in the thousands). Consequently, prestoring all permutations for selection is not feasible. Additionally, because each of the many thousands of traded securities provides a different risk profile and is traded in discrete units, the space of permutations is both discrete and explosive. Hence, using a straightforward generate-and-test approach (e.g., with conventional mathematical programming techniques) is unlikely to work well.

To constrain the combinatorial nature of this search problem, we can abstract all elementary piecewise linear functions having a similar shape into one qualitative function with linear edges having a slope of 1, 0, or -1 over qualitative ranges on the real-line. For example, all functions having one edge with slope 0 over $(0, x_i) \in \Re$ and another edge with a positive slope a_i over $(x_i, \infty) \in \Re$, where a_i and x_i differ across these functions, are replaced by one qualitative function having one flat edge over $(0, x)$ and another edge with slope 1 over (x, ∞), where $x \in (0, \infty)$ is a qualitative point. This abstraction lowers drastically the number of elementary functions, rendering the use of a simple generate-and-test approach computationally feasible.

At the same time, this qualitative abstraction also results in the loss of important information. One implication is apparent in the case of scenario 3. The butterfly goal function is synthesized using two similar elementary functions, corresponding to the purchase of two call options, one with strike price s_L and another with strike price s_H. Because these two functions are now represented by the same abstract function, this goal function cannot be synthesized unless we can rediscover lost information about the ordering of strike prices. Hence, we need to use heuristic search operators capable of rediscovering lost information by stretching and steepening edges in permutations of abstract elementary functions.

Qualitative Synthesis

Qualitative synthesis (QSYN) is a QR technique that can solve this synthesis problem (20). The systemic concepts underlying QSYN are as follows.

1. A security is a two-terminal component (system) whose input node is some risky economic parameter and output node is its value.
2. A risk profile describes the behavior of a component (system) over all its operational regions. It describes the output contingent on the input only at the end of some risky period. Hence, unlike QSIM, QSYN assumes that the input and output nodes are time insensitive.
3. An abstract risk profile describes the qualitative behavior of all components of the same type.
4. The qualitative behavior of a system is the sequence of qualitative states that the output node exhibits as the input node varies over the entire range of values it can take on. However, while in QSIM $qdir \in \{1, 0, -1\}$, in QSYN $qdir \in \mathbb{N}^+$.
5. Because the risk profile of a position is the algebraic sum of risk profiles of its security components, a position is a two-terminal system made from components connected in parallel. That is, where the behaviors of a system and its components are analogized to *transfer functions* (31), given the transfer functions of any two components, their sum is the transfer function of a sys-

tem made from the two components connected in parallel.

Based on these concepts, the problem is one of synthesizing the structure of two-terminal systems—identifying sets of structurally connected components—that produce some goal qualitative behavior. QSYN solves this problem using the following basic search approach. It takes as input: the qualitative behavior G of a prospective system (e.g., butterfly risk profile), and a set $\mathcal{Q}$ of n qualitative behaviors $Q_1, Q_2, \ldots, Q_n$ that each abstracts the qualitative behaviors of all components of the same type. Upon selecting a pair of behaviors in $\mathcal{Q}$, Q_i and Q_j ($i \neq j$), a permutation is created as their sum. This permutation is then compared against G. If it matches part of G, it is added to $\mathcal{Q}$ with a reference to the Q_i and Q_j used to create it. If it matches all of G, a parallel connection of components i and j is identified as one possible way to synthesize the prospective system. These steps are repeated for every possible permutations involving a pair of different behaviors in $\mathcal{Q}$, including pairs containing partially matching permutations newly added to $\mathcal{Q}$. In so doing, QSYN finds all permutations of elementary behaviors in $\mathcal{Q}$ that match G.

QSYN uses two means to deal with flaws in its basic search approach. First, to avoid an exhaustive search of the space of combinations of behaviors in $\mathcal{Q}$, QSYN uses a goal-directed search process. Because the number of possible permutations is on the order of $\|\mathcal{Q}\|^2$, QSYN constrains the generation of permutations using knowledge about the additivity of qualitative behaviors (implied by the transition rules QSIM employs). For example, if the *qdir* in both the first intersecting states in a permutation of two behaviors is 1 (increasing) and the *qdir* in the first state in the goal behavior is 0, QSYN readily prunes that permutation because it will not yield a match. Hence, instead of computing the sum of two behaviors in $\mathcal{Q}$ and then comparing it against the goal, QSYN compares the goal behavior to the sum of the two combined behaviors as this sum is being computed gradually, one pair of states at a time. The notions of *sum* and *match* are defined as follows.

- For the *sum* of two behaviors, denoted $\oplus$, consider a behavior to be a sequence of elements of the form [(IN $\langle qdir\ qval\rangle$)(OUT $\langle qdir\ qval\rangle$)]. Furthermore, assume the existence of behaviors Q_1 and Q_2, with m and n elements, respectively, and let [ ] denote the kth element in a behavior. Elements $Q_1[i]$ ($1 \leq i \leq m$) and $Q_2[j]$ ($1 \leq j \leq n$) are *corresponding*, if the IN-*qval* of $Q_1[i]$ is contained in the IN-*qval* of $Q_2[j]$, or vice versa. The sum of two elements, $Q_1[i] \oplus Q_2[j]$, is a new element, $Q_3[k]$, in which: (1) IN-*qval* is the intersection of IN-*qvals* of $Q_1[i]$ and $Q_2[j]$, and (2) OUT-*qdir* is the algebraic sum of OUT-*qdirs* of $Q_1[i]$ and $Q_2[j]$. For example, assuming that $(x_1, x_2) \subseteq (x_1, x_3)$:

$$\begin{array}{ll} Q_1[i] & = [(\mathrm{IN}\langle *\ (x_1, x_2)\rangle)(\mathrm{OUT}\langle\ \ 1\ *\rangle)] \\ Q_2[j] & = [(\mathrm{IN}\langle *\ (x_1, x_3)\rangle)(\mathrm{OUT}\langle -1\ *\rangle)] \\ \hline Q_1[i] \oplus Q_2[j] & = [(\mathrm{IN}\langle *\ (x_1, x_2)\rangle)(\mathrm{OUT}\langle\ \ 0\ *\rangle)] \end{array}$$

The sum of two behaviors, $Q_1 \oplus Q_2$, is thus the sum of every pair of their corresponding elements.

- For the notion of *match*, two corresponding elements are matching, denoted $Q_1[i] \approx Q_2[j]$, if they have the same OUT-*qdir*. For example, although the two preceding sample elements are corresponding because $(x_1, x_2) \subseteq (x_1, x_3)$, they do not match because their OUT-*qdirs* are 1 and -1. Two behaviors are matching, $Q_1 \approx Q_2$, if every pair of their corresponding elements match. A behavior Q_1 partially matches another behavior Q_2, if Q_1 matches the first few consecutive elements of Q_2.

QSYN deals with another flaw in its basic search approach that relates to the loss of important information due to the fact that each qualitative behavior in $\mathcal{Q}$ abstracts all the behaviors of components of the same type. QSYN uses two heuristic synthesis operators—STRETCH and STEEPEN—on elements of the behaviors in a permutation, to rediscover the information lost. It is easiest to understand how these operators work by looking at the next example.

Applying Qualitative Synthesis

Let's go back to the problem in scenario 3, which entails the synthesis of a butterfly risk profile. One of the permutations of risk profiles that QSYN tries includes the pair of elementary risk profiles denoted Q_i and Q_j at the top of Fig. 5. Apparent from Fig. 4, at least part of G can be synthesized using Q_i and Q_j. Yet, $Q_i \oplus Q_j \neq G$ because Q_i and Q_j are each an abstraction of an entire class of risk profiles with the same qualitative shape. QSYN, hence, tries to use operators STRETCH and STEEPEN so as to synthesize G using Q_i and Q_j.

The next discussion traces QSYN's synthesis process, corresponding to the emphasized path in the search tree shown in Fig. 5. Starting with the first triplet of elements at the top of the tree, QSYN concludes that $Q_i[1] \oplus Q_j[1] \approx G[1]$. For the next triplet, it concludes that $Q_i[2] \oplus Q_j[2] \neq G[2]$ because the OUT-*qdir* of $G[2]$ differs from the OUT-*qdir* of $Q_i[2] \oplus Q_j[2]$. But, because the OUT-*qdir* of $G[2]$ is equal to the OUT-*qdir* of $Q_i[2] \oplus Q_j[1]$, a modified version of Q_j, denoted Q_j' in Fig. 5, in which the first element is stretched over the IN-*qual* $(0, s_M)$, is more likely to contribute to the synthesis of G. QSYN, hence, uses operator STRETCH to extend $Q_j[1]$ over $(0, s_M)$ and to conclude that $Q_i[2] \oplus Q_j'[1] \approx G[2]$. For the next triplet of elements QSYN concludes that $Q_i[2] \oplus Q_j'[2] \neq G[3]$, because the OUT-*qdir* of $G[3]$ differs from the OUT-*qdir* of $Q_i[2] \oplus Q_j'[2]$. However, this mismatch can be eliminated by modifying the OUT-*qdir* of $Q_j'[2]$ from -1 to -2. QSYN therefore applies operator STEEPEN to create a new version of Q_j', denoted Q_j'' in Fig. 5, and to conclude that $Q_i[2] \oplus Q_j''[2] \approx G[3]$. At this point QSYN found a partial match. It hence adds $Q_i \oplus Q_j''$ to $\mathcal{Q}$ as a new "elementary" behavior and then continues to synthesize G in the same fashion.

How can we interpret the partially matching permutation QSYN synthesized? This permutation is made from two elementary risk profiles that were modified by operators STRETCH and STEEPEN. These modified risk profiles provide information about how to configure a position whose risk profile partially matches a butterfly. First, Q_i and Q_j'' have the qualitative shape of the risk profiles of a buy call option position and a sell call option position, respectively. Second, because $s_L \leq s_M$, the strike price of the purchased call s_L must be smaller than that of the sold call s_M. Last, the absolute

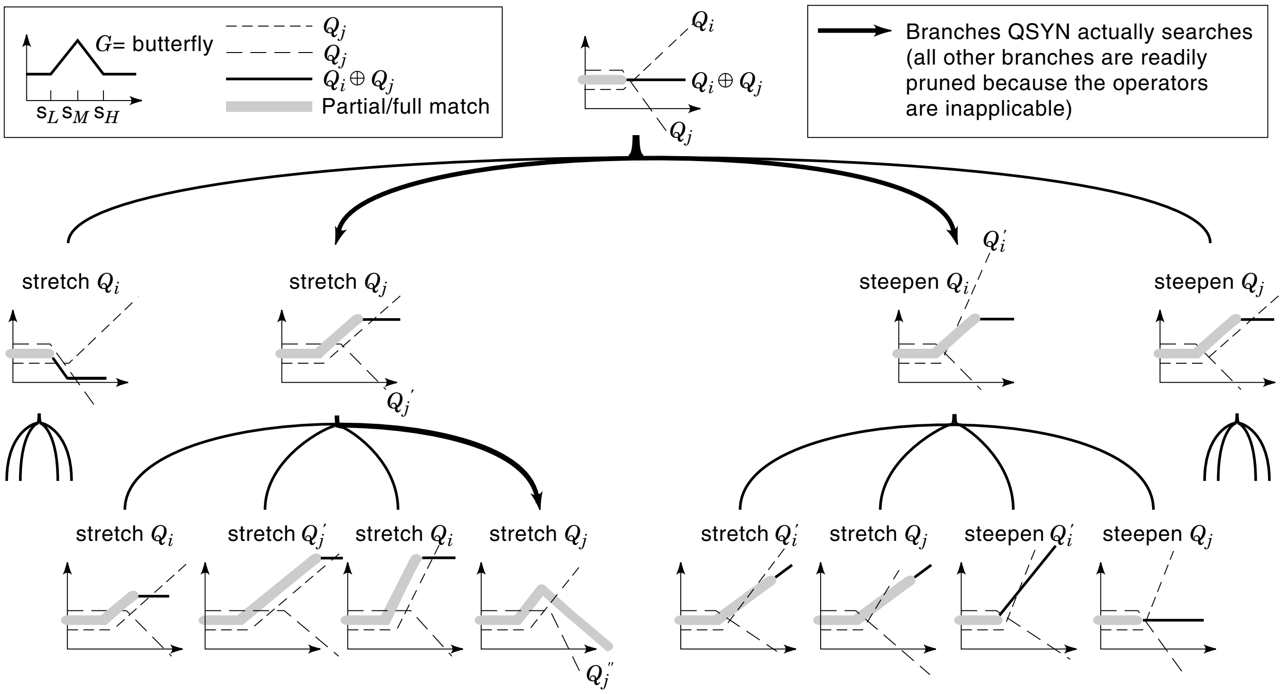

Figure 5. QSYN synthesizes part of a butterfly risk profile. Apparent from the tree encompassing only a small part of QSYN's search space, QSYN vigorously prunes the search space using heuristic search operators *stretch* and *steepen.*

value of the OUT-*qdir* of the second element in Q_j'' is 2, indicating the need to sell more than one call for every purchased call. This information is identical to the one provided in the solution to scenario 3.

CONCLUSION

This article focused on the strategic role of QR techniques in investment decision making. QR techniques can support and augment the way financial decision makers reason about systems involving a high degree of *internal* uncertainty. Even though external uncertainty is a result of uncontrollable factors in the environment, internal uncertainty stems from many complex interdependencies between parameters that must be understood in order to control the behavior of a system. Quantitative financial models aim at providing insights that can reduce the internal uncertainty. But, because of their limitations, reasoning with these models qualitatively can help to reduce internal uncertainty further. In this sense, QR techniques help to leverage the use of quantitative models, especially in early decision-making stages. In these stages, a qualitative understanding of a problem is vital. It determines, however implicitly, the alternative courses of action evaluated and compared at later stages through a detailed quantitative analysis.

In light of the relative success of QR techniques, recent research seeks to expand the scope to which these techniques can be applied to financial and economics problems. This research includes attempts to expand standard QR techniques (32), to enable their use with a broader range of models involving, for example, an average parameter (X/Y) that "tracks" a marginal parameter (dX/dY); choices that economic agents make to select between alternative courses of action (behaviors); and uncertain parameters modeled in mean-variance terms, using continuous probability distributions.

Relatedly, other research stresses the need to adapt existing QR techniques to complex economic and financial problems (33). Because standard QR techniques are typically developed for physics and engineering problems, they ignore the significant contribution that mathematical economics has made to the study of dynamics (e.g., in the area of stability) through exploitation of the idiosyncratic structure of economic systems. This work also points out possible implications of using QR techniques on related areas such as econometrics.

Focusing on key concepts that make QR techniques appealing, other recent work applies these concepts using other artificial intelligence technologies such as fuzzy logic (34).

BIBLIOGRAPHY

1. M. J. Lyneis, *Corporate Planning and Policy Analysis: A System Dynamics Approach,* Cambridge, MA: MIT Press, 1980.
2. E. Helfert, *Techniques of Financial Analysis,* Homewood, IL: Richard D. Irwin, 1982.
3. R. A. Brealey and S. C. Myers, *Principles of Corporate Finance,* 3rd ed., New York: McGraw-Hill, 1988.
4. H. M. Blalock, Jr., *Causal Inferences in Non-experimental Research,* Chapel Hill, NC: University of North Carolina Press, 1964.
5. S. A. Zenios, *Financial Optimization,* Cambridge, MA: Cambridge University Press, 1993.
6. M. J. McInnes and W. J. Carlton, Theory, models and implementation in financial management, *Management Science,* **28** (9): 957–980, 1982.

7. M. J. Bouwman, Human diagnostic reasoning by computer: an illustration from financial analysis, *Management Science,* **29** (6): 653–672, 1983.

8. J. M. Bouwman, P. A. Frishkoof, and P. Frishkoof, How do financial analysts make decisions? A process model of the investment screening decision, *Accounting Organizations and Society,* **12** (1): 1–29, 1987.

9. A. Tversky, D. Kahneman, and P. Slovic, *Judgment Under Uncertainty: Heuristics and Biases,* Cambridge, MA: Cambridge University Press, 1982.

10. H. A. Simon, Theories of bounded rationality. In C. B. McGuire and R. Radner (eds.), *Decision and Organizations,* Minneapolis, MN: University of Minnesota Press, 1986, pp. 161–176.

11. G. F. Davis, Agents without principles? The spread of the poison pill through the intercorporate network, *Administrative Science Quart.,* **36** (4): 583–613, 1991.

12. R. Berndsen and H. Daniels, Causal reasoning and explanation in dynamic economic systems, *J. Economic Dynamics and Control,* **18** (3): 251–271, 1994.

13. M. Benaroch and V. Dhar, Controlling the complexity of investment decisions using qualitative reasoning techniques, *Decision Support Systems,* **15** (2): 115–131, 1995.

14. G. J. Wyatt, R. R. Leitch, and A. D. Steele, Qualitative and quantitative simulation of interacting markets, *Decision Support Systems,* **15** (2): 105–113, 1995.

15. W. Hamscher, Model-based financial data interpretation, *Proc. of the First International Conference on AI Application on Wall Street,* pp. 178–187. New York: IEEE Press, 1991.

16. P. Alpar and W. Dilger, Market share analysis and prognosis using qualitative reasoning, *Decision Support Systems,* **15** (2): 133–146, 1995.

17. C. Apte and S. June-Hong, Using qualitative reasoning to understand financial arithmetic, *Proc. AAAI—87,* pp. 942–947. AAAI Press, 1987.

18. R. Berndsen, Causal ordering in economic models, *Decision Support Systems,* **15** (2): 157–165, 1995.

19. L. Kuan-Pin and A. M. Farley, Causal reasoning in econometric models, *Decision Support Systems,* **15** (2): 167–177, 1995.

20. M. Benaroch, A technique for qualitatively synthesizing investment vehicles, *Information Sciences,* **90** (1–4): 281–302, 1996.

21. S. Raghunathan, An application of qualitative reasoning to derive behavior from structure of quantitative models, *J. Management Information Systems,* **11** (1): 73–100, 1994.

22. F. Gardin, R. Power, and E. Martinelli, Liquidity management with fuzzy qualitative constraints, *Decision Support Systems,* **15** (2): 147–156, 1995.

23. J. E. Elton and J. M. Gruber, *Modern Portfolio Theory and Investment Analysis,* 5th ed., New York: John Wiley & Sons, 1995.

24. D. T. Coggin, Active Equity Management, pp. 51–72. In F. J. Fabozzi (ed.), *Portfolio and Investment Management,* Chicago: Probus Publishing Company, 1989.

25. B. Kuipers, Qualitative simulation: Then and now, *Artificial Intelligence,* **59**: 133–140, 1993.

26. B. Kuipers, Qualitative simulation, *Artificial Intelligence,* **29**: 289–338, 1986.

27. D. Gentner and A. S. Stevens, *Mental Models,* Hillsdale, NJ: Erbaum Publishers, 1983.

28. J. P. Hayes and S. P. Weinstein, Adding value to financial news by computer, Proc. First International Conference on Artificial Intelligence Applications on Wall Street, pp. 2–8. New York: IEEE Press, 1991.

29. C. Smith, C. Smithson, and D. Wilford, *Managing Financial Risk,* New York: Harper Business Publishing, 1990.

30. J. Hull, *Options, Futures, & Other Derivative Securities,* Englewood Cliffs, NJ: Prentice Hall, 1993.

31. R. Saucedo and E. E. Schiring, *Introduction to Continuous and Digital Control Systems,* New York: Macmillan, 1968.

32. M. Benaroch, Extending qualitative reasoning techniques for use with advanced dynamic economics and financial models. Working paper, School of Management, Syracuse University, April 1997.

33. M. Benaroch, Artificial intelligence in economics: truth and dare, *J. Economic Dynamics and Control,* **20** (4): 601–606, 1996.

34. V. Deinichenko, G. Bikesheva, and A. Borisov, Fuzzy approach in economic modelling of economics of growth, in P. Ein-Dor (ed.), *Artificial Intelligence in Economics and Management,* pp. 155–173. Norwell, MA: Kluwer Academic Publishers, 1996.

MICHEL BENAROCH
Syracuse University