

Personalizing Online Trading Without Spoiling Its Intrinsic Advantages

Myron Sheu, (Email: mshou@csudh.edu), California State University, Dominguez Hills
Stewart Shen, Old Dominion University

ABSTRACT

Although access to the Internet has become more efficient than before, investors seem still shy away from taking advantage of online trading. Early research indicates that one main reason for the setback is the impersonality of online trading. Brokerage firms and researchers have investigated various kinds of decision support that mimic offline brokerage services, but the lack of timeliness and simplicity of such support may have just ruined the purpose of online trading. To cope with the situation, we have investigated a heuristic model for representing financial securities such that inexperienced investors can naturally perceive personalized guidance for their investment decisions. In addition to a description of our model, this paper also discusses the reasoning schema behind the model and proposes the methods for constructing the model. In contrast to other approaches, our model is aimed to provide personalized decision support without spoiling the intrinsic advantages of online trading.

NEEDS FOR PERSONALIZED ONLINE INVESTMENT SERVICES

As millions of less sophisticated investors turn to active online trading, we believe that the ultimate success in shifting the paradigm of investment from offline to online depends on investor confidence in the online trading environment. Following initial enthusiasm in online trading, many investors have started considering the current online trading environment not only impersonal but also deceptive (Barber & Odean, 2001). To cope with the situation, online brokerage firms have rushed to provide decision support in a bid to change the trade-by-rumor atmosphere of online investing (Matthews, 2004). Until now, decision support for online trading has been provided through supplementary educational systems and stock-analysis systems in a way that is, unfortunately, only suitable for analysts (Larkey, 2002). Two better-known examples, i.e., the Foundational and Signature Advice System (FSAS) from Charles Schwab (2004) and Holding Company Depository Receipts (HOLDR) from Merrill Lynch (2004), have made noteworthy progresses in providing investors with effective decision support. However, they did so at the cost of giving up some advantages of online trading. Specifically, FSAS offered in-depth analyses of individual stocks and gave recommendations in terms of personal risk tolerance in traditional settings, that is, by phone conversation or by meeting in person. HOLDRS allowed individual investors to trade stocks within a sector or industry in a bundle or in individual pieces. By doing so, HOLDRS discouraged excessive holdings of specific stocks while offering flexibility for adjusting underlying stocks in accordance with an individual's financial goals. However, the puzzle on how to adjust was still upon each individual. Recently, Mertens, Marius, and Yuliadi (2004) proposed a rule-based expert system to help online decision-making through an intelligent dialog model. The key design of this dialog model is that the system reacts intelligently to users' input. Due to the interactive reasoning process, such a decision support model demands appropriate trainings and strict disciplines on the user side while it incurs responsibility that may be hard to bear on its own side. Understandably, the complexity of such an approach to decision support also challenges the affordability of most online trading firms, especially small and medium-sized ones. As financial markets increasingly operate over the Internet, a wide variety of data sources have become readily available to ordinary investors. In such efficient markets, in our view, the focus of online brokerage services should not be on providing users with an overwhelming amount of data. Instead, it should be on helping investors quickly comprehend the markets by providing only relevant data in a personal manner as being done in the offline trading environment (Kauffman, Subramani, & Wood, 2004).

In an attempt to personalize online decision support, our research has been exploring suitable individualized knowledge representation models that are easily comprehensible to online investors. By focusing on knowledge representation of the fundamental valuation of an investment security, we help online investors make their own wise, informed decisions on investment without poisoning the character of online trading. We agree with Rajgopal, Venkatachalam, and Kotha (2000) that being able to articulate the type and degree of an investment risk to an online investor in accordance with his/her financial goals would benefit both sides of online trading. This belief has guided our efforts on developing our format of decision support. Our knowledge representation model intends to facilitate investors' assessment of investment risks in the context of each individual's own risk tolerance against the corresponding opportunities. The main challenges to our approach are that financial markets are extremely dynamic and individuals have unique and changing financial situations and goals. Accordingly, we think three important quality features of the online investment environment must be upheld. The first quality feature to address is the ability for an online investment environment to understand each individual's investment goals and the economic potential of a company or an industry, both of which may evolve over time. The second key feature to address is the ability to adapt to the change of an individual's investment goal. Finally, the third vital attribute is to furnish the first two features in an efficacious manner.

Our approach is intended to match some mental processes that help human beings organize their thoughts (Formica & Missikoff, 2004). Challenges to attaining these advantages do exist, however. Individual financial situations and investment goals are extremely diverse and ever changing and thus they demand personalized, dynamic financial advice. Hence, any tailored knowledge representation of financial services must be highly adaptable. To achieve the adaptation of this kind, we believe that the online investment environment should be constructed significantly differently from current ones and needs to incorporate additional mechanisms for recognizing, justifying, and representing exceptional valuations of an investment opportunity in light of exceptional characteristics of an individual's financial goals. In the rest of the paper, we discuss our model for delivering these features.

RECOGNITION MECHANISMS FOR EXCEPTIONAL VALUATIONS

Our knowledge representation model for online trading is in an object-oriented approach, which is considered as a primary framework for semantic ontology (Kim, 2002). In response to the dynamic nature of investing, our model first recognizes an individual's exceptional valuations that contradict to a default valuations resulting from inheritance (Reiter, 1980; Reiter, 1985). To collect exceptional observations, we embed some data collection methods into the model so that exceptional valuations can be observed long before they grow to possess consequential significance. Additionally, the model integrates multiple reasoning schemas so that a default reasoning mechanism measures the market risk and an explicit reasoning mechanism recognizes an individual's unique risk. Without exceptional consideration, the unique risk of an individual security would be overlooked and then it would be left to his/her unaided speculations. Since the significance of the unique risk could vary more widely than the market risk, an early, accurate detection of the unique risk would play a critical role in presenting an investment opportunity to an online investor in a concrete and measurable form.

In general, exceptional valuations can be recognized in two reasoning approaches, namely, deduction and induction (Genesereth & Nilsson, 1988). In deductive reasoning, an individual is able to recognize the exceptions that have been cultivated by numerous samplings that are collected and analyzed at a central place. An example of exception resulting from deductive reasoning is what follows. If a trading pattern indicates that traders conduct more buying trades on Mondays even if the market does not reach low and conduct more selling trades on Fridays even if the market does not reach high. An exceptional investment opportunity is then detected when the market generally moves up and down in an opposite direction on these two respective weekdays. Note that such an exception is established initially through induction by observing many cases and is then through deduction by verifying with additional cases. Such rules are provided to the clients as default exception rules. Sophisticated investors may choose to accept such rules as given, or they may choose to override them with the refined rules based on own experiences and preferences. Less sophisticated traders may accept it or withhold it while learning from it. If so, the less sophisticated individuals should be reminded of the consequences of their decisions.

Recognitions of exceptions may also be originated from specific performances of a particular security. For example, certain stocks are observed to start moving up soon before reporting quarterly results, and then drop sharply immediately after the reporting. After learning through induction and deduction, such rules associated with specific stocks are returned to the clients, who can again inherit, override, withhold, or reject them.

Some seemingly highly valuable exceptions may be learned through induction, but they may take a lengthy period to establish the confidence through deductive applications. After repeatedly observing the occurrences contradictory to default results, consistent findings are organized and identified as additional exceptions. They may be valuable to the clients and should be properly presented to them. But inductive reasoning usually does not possess the level of confidence that deductive reasoning offers. Thus such exceptions established by inductive reasoning need to be clearly marked with appropriate levels of confidence, which are often relatively weak (Dalkey, 1988), before given to the clients.

Another kind of exceptions that are appropriate for inductive reasoning may be related to some particular investment performance or financial decisions of an individual. For example, because of limited access to information, one may primarily invest exclusively in a few large companies or only in domestic companies. To help improve the long-term health of the individual's portfolio, an exception should be identified as a reminder to the individual that he should restrain from investing further heavily in these stocks and the like. Note that when a new exception has just been created its associated confidence level tends to be weak, but it gradually strengthens if more observations are supportive of the exception, or it may further weaken.

Along with a degree of confidence, exception recognition should also possess a degree of granularity to indicate that some are applicable to individual securities while others are applicable to one's entire portfolio. An example of granular exception may be that one's entire portfolio should move away from certain sectors during down cycles of them or the trading frequency should be reduced if the performance of one's portfolio is worse than some indexes. The resultant investment strategy may include a default rule and the two exceptional considerations, as listed in the following:

- Default Trading Pattern: Follow the model portfolio established upon one's financial goals.
- Exception Trading Pattern I: Reduce trading frequency.
- Exception Trading Pattern II: Reduce investment in certain sectors (e.g., telecommunications).

The above two exceptions both have a high granularity and are not applicable to any individual security. If no exceptions were raised, one should manage his/her investment portfolio by following general financial equations amid personal risk preference. That is, without particular reasons, the influence to an investment decision should be driven by one principle, namely, to follow a model portfolio that has been configured in terms of the individual's long-term financial goal.

In addition to exceptions applicable to one's entire portfolio, other exceptions may be associated with individual securities. These exceptions for individual securities usually reveal unique characteristics of particular securities. A security that performs in a way contradictory to the majority of securities in the industry in which it belongs could be recognized as an exception. As an example, the market values of some stocks, reflecting the cyclicity of underlying business, change significantly during certain economic times. Acquiring these stocks for a short term at the right times may often yield a very good result. Identifying these exceptional periods may be desirable to someone with more risk tolerance. In an object-oriented paradigm, by inheritance, a security may not be typically recommendable as a long-term investment but, by polymorphism, a short-term investment opportunity is recognized when the security hits a cyclic bottom. In fact, exceptional valuations should reflect most short-term buying opportunities while the general representational model fits one's long-term goal.

On the other hand, exceptions could also be used to indicate unnecessary and unique risks (Sharpe, Alexander, & Bailey, 1999). For example, because excessive concentration on specific industries may somewhat deteriorate the quality of one's investment portfolio, an exception could be attached to selected securities that together

constitute a heavy concentration in a few industries. As a result, individual investors are aware of such exceptional risks and should balance the distribution of their investment baskets accordingly.

The above discussion attempts to characterize exceptions that should be furnished to individual online investors. Exceptions may be represented at the portfolio level as well as at the individual security level. At either level, exceptions could arise from different perspectives; some may depict significant unique risk while others may reflect exceptional values. A personalized representation of these exceptions could significantly compensate the impersonality of online trading and thus help millions of unsophisticated cyber investors make sound financial decisions.

REPRESENTATION SCHEMA FOR EXCEPTIONAL VALUATIONS

While representing exceptions was considered in the object-oriented paradigm (Minsky, 1975), the fact that exceptional valuations dynamically change in light of both qualitative and quantitative significance imposes great challenges to tailoring such representations. In the financial market, the exceptional valuations of a security often result from a variety of factors. Each of these factors may in turn result from some other reasons. Exhaustively expressing all causes may be confusing for a trader to understand and oftentimes may be infeasible to implement. Besides, not all causes are equally influential to the long-term quality of a security. Furthermore, the significance of the influencing effect of an exception may vary over time, and thus its varying influences to a security should be accurately represented as they change. In the rest of this section, we exemplify the schema developed in our research for organizing and representing exceptional factors.

Multiple Inheritances

Multiple factors may affect the exceptional valuation of a security. Default reasoning alone cannot adequately accommodate exceptional factors since such factors can quickly gain or lose influences to the valuation of a security investment while exception reasoning may also come and go. Therefore, for tailored decision support, our model of knowledge representation extended the single-parent inheritance to the multiple-parent inheritance as shown graphically in Figure 1. Whenever all the unique risks taken by an investor become insignificant, the representation model for the investor could reduce back to the traditional hierarchical model. However, in another extreme situation in which the exceptional factors for an investor should overwhelmingly dominate his/her investment strategies, the weight of the default hierarchy could diminish from the multiple-inheritance hierarchy and, as a result, the unique risk tolerance could override a default valuation. Our model thus capably cultivates unique investment strategies as the investor's unique financial characteristics evolve.




Figure 1. Extra risk tolerance may be justified if exceptional characteristics of an investor are taken into account.

In a multiple-inheritance model, a bold investment decision may conflict with the risk constraint from one or more parent objects but does not necessarily trigger any violation to the investment principles of an individual so long as the composite portfolio at the parent level still remains efficient.

Composite Object Class

Due to the use of multiple inheritances, the child object class often becomes too complex for a user to observe the details. For a better alternative, we introduced a composite object class. A composite object aggregates multiple objects, each of which could inherit from a different parent object class. With the introduction of composite object class, a multiple-inheritance hierarchy can also be viewed as a composition of multiple single-inheritance hierarchies. Figure 2 shows such a flexible representation that facilitates visualizing the synthesized and featured risk tolerance.




Figure 2. A synthesis of single inheritance facilitates visualizing the synthesized risk tolerance.

In our representation model, the knowledge derived from statistical analysis is displayed using generalization and aggregation. Our intent was motivated by the fact that statistical expressions were often too complicated to express the subtle differences among highly diverse investors, and they are too procedural to adjust in accordance with the changing financial situation of an individual. As a result, they are too obscure to be understood by a naïve investor, especially under time pressure in online trading. Figure 3 shows that, supported by a multi-inheritance hierarchy, composite objects can show synthesized efficiency, as depicted under the first indifferent curve, and can show segmented efficiency, as depicted under the second indifferent curve.




Figure 3. Composite objects in a multiple inheritance hierarchy support the aspect analysis as they together represent a changing investment goal.

While the model offers an overview of the uniqueness of an investment portfolio, it appears to be highly declarative and in fact has to rely on additional mechanisms in order to be structural and evolvable. The methods for measuring unique risks, to be described in the next section, will reveal associations between the exceptional valuations of a security investment and the exceptional characteristics of an individual’s financial goals.

JUSTIFICATION METHODS FOR EXCEPTIONAL VALUATIONS

Our model associates an exception with a measuring mechanism to justify the confidence of the exception. Such a mechanism is required for the ability to learn in an unsupervised manner (Rich & Knight, 1992) and must be intuitive because of the impersonality of online trading. In choosing an appropriate justification method, we have assumed that the financial market changes randomly and does not make movements in response to typical individual actions, unless a large block transaction is requested. Another assumption of our investment model is that the return by taking a market risk is estimated by using the average return of a suitable market such as the NASDAQ index.

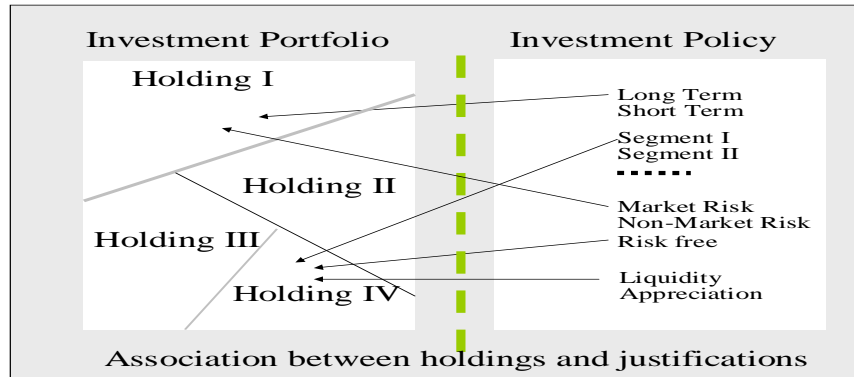


Figure 4. A binding structure that ties an investment portfolio to tailored investment goals.

Several logical reasoning schemas can provide the justification of an exceptional valuation, but we have found leveraging a decision tree to be most appropriate due to its simplicity and intuitiveness. A decision tree, which can also be displayed as a decision table, can be dynamically generated in terms of an individual’s trading records and financial profile. The decision nodes on the tree exhibit exceptional factors while the decision branches indicate the likelihood of each exceptional factor. The decision tree should underpin the knowledge representation hierarchy for a security or a portfolio described in the previous section. As discussed earlier, the exceptional valuations of a stock are perceived in accordance with personal financial profiles. As shown in Figure 4, such correspondences can be maintained within an individual profile. The most critical issue we have to deal with is that as the decision nodes increase, the required amount of computation exponentially increases and thus may threaten the feasibility of such a reasoning structure for online trading. As a remedy, pruning a decision tree is necessary to reduce the amount of time and space required in support of decision making. The fact that the relevance of decision nodes rely on conditions which often do not occur simultaneously offers substantial savings through pruning. Specifically, we construct each decision tree by grouping only the factors relevant at that time to an individual profile and prune the rest. For example, if a company appears quite typical in its native industry except for its exceptionally strong or weak growth rate, then the relevant decision tree should only include decision nodes that depict the strength of growth.

Table 1: A Two-Tier Decision Tree For Justifying Exceptional Valuation Of A Security

Tier	Factor	Possibility			
		1	Growth	Faster than the industry: 70%	
2	Market share	Increase: 65%	Decrease: 35%	Increase: 65%	Decrease: 35%
Leaf Node	Appreciation of stock value	45.5% Higher Appreciation	44% (=24.5% + 19.5%) Average Appreciation		10.5% Lower Appreciation

The significance of each factor varies according to the changing market and business conditions. Consequently, some factors weigh much heavier than others do at different times and the most significant ones are chosen to constitute a decision tree for validating exceptional valuations. For example, if the growth rate and market share were determined to be the only significant factors to the exceptional valuation of a security, then there could be a decision tree, as exhibited in Table 1, for justifying its exceptional valuation:

- The first tier indicates the growth rate with two branches to indicate a higher-than-average rate with a probability of 0.7, and a lower-than-average rate with a probability of 0.3. Thus, the materialization of exceptional valuation for growth rate may statistically be 40%, which results from 0.7 minus 0.3.
- The second tier exhibits the market share rate. The chance of having an expansion on market share is 30%, which results from 65%, the chance of increased market share, minus 35%, the chance of decreased market share.
- The leaf nodes indicate the combined possibilities of two tiers. The chance of materializing exceptional valuation is around 45.5%, resulting from 70% times 65%. Subsequently, the chance of having an average appreciation is around 44%, resulting from 70% times 35% plus 30% times 65%. Finally, the chance of having lower than average return is 10.5%, resulting from 30% times 35%.

Nevertheless, the historical data might suggest some correlation between the growth rate and the market share, the estimated range of probability for exceptional valuation could shift significantly higher, from 45.5%, resulting from .7 times .65 when the two factors are independent, to around 67.5%, resulting from $(.7 + .65) / 2$ when they are highly correlated. The degree of correlation could be determined by examining the historical data so that the decision tree can incorporate such correlations. Similarly, the decision tree can display the chance of having a lower return as a range from 10.5% to 30%.

Note that default valuations are applied in general and specific exceptional valuations are applied only when they are applicable to specific investment segments. Figure 5 demonstrates a general decision tree for a stock that may possess some unique valuation. Accordingly, only the portion of the tree reflects the exceptional valuation is calculated and the rest are pruned. The measure not only makes the model computationally feasible but also restrains it from making radical recommendations.




Figure 5. The attributes that only yield the possible market valuation are pruned for their irrelevance to justifying the exceptional valuation.

At the bottom of the decision tree, there are leaf nodes identifying worst possible and best possible results and the distribution within the range should normally fall into a bell shape. For example, a security might appreciate its valuation at the worst possibility of 45.5%, resulting from 0.7×0.65 , or at the best possibility of 67.5%, resulting from $(0.7 + 0.65) / 2$. Again, if the data of correlation between involved attributes are available, the significant range of chance could be much narrow because of the improved worst situation. At the next level, the exceptional valuation from each decision trees could be summed up along with the default valuations. Although it is going to be complex if multiple exceptions are considered, most often one or two indicators of exception suffice for estimating the total exceptional valuation, which is similar to the cost drivers in cost accounting and can be easily determined through sensitive tests. While in general decisions trees should be hidden from visual exhibition, it is possible to be visible in

case an investor is curious about the reason beings behind the exceptional valuation. Since the display of a decision tree usually involves extensive graphical items, we think decision tables should result in better performance since theoretically decision trees and decision tables are interchangeable. However, we admit that the hierarchical display of a decision tree appears much friendlier than a tabular structure that does not clearly show the logical dependency, for which we will offer some suggestions when we discuss implementation issues in the next section.

The above examples demonstrates how the decision-tree approach to reasoning could convince an individual not to make common mistakes, such as buying high and selling low. In addition, it can also be applied to preventing an individual from repeating his/her unique mistakes in light of the individual's unique investment objectives. The exceptions of this kind may be specific trading behavior that leads to a mismatch between one's investment goal and investment practice. The emphasis is then changed to the uniqueness of their mistakes rather than the common mistakes induced by the financial market.

IMPLEMENTATION ISSUES

While the proposed approach to providing online investors with tailored decision support may sound appealing, the questions of feasibility may emerge. Primarily, one may ask these questions: Is it possible for a brokerage firm to support millions of custom portfolio analysis? Is it possible for an individual to easily assemble the information closely related to his/her financial portfolio over the Internet? In short, our answer is quite positive largely because many new features that are supportive of object-oriented applications have been integrated into a relational database system, including inheritance, user-defined data types, composite data types, and a variety of function-specific objects.

While the technology is available for implementation, the viability lies in the architecture of an online trading environment that thus far is no match for an offline trading environment. Imagine how an investor is served when he/she walks into a brokerage office. Specifically, an offline broker may carefully control the information flow presented to the investor and heuristically convince what should be the right kind of investment for each individual. To gain trust of this kind, which is vital to winning the trading business, an online brokerage firm should completely redesign its trading environment to emulate the offline investment environment, including not only the visual layout but also the management of information flows. Figure 6 illustrates such a model for online trading developed in our research. The central component of the model is a UI management layer for selectively applying investment principles embedded in domain knowledge objects for implicit reasoning. These components are invisible and merely for controlling the information flow to an online investor. In a variety of granularity they are the major intelligent entities behind a visual UI.

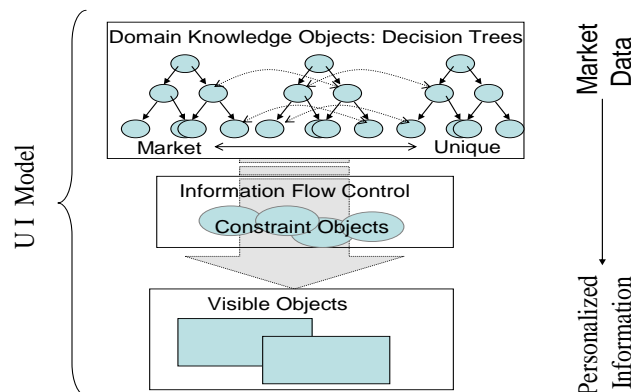


Figure 6. An exemplification of the UI model, which manages both the visible and invisible components of online trading environment.

Additionally, constraint objects behind the scene could be dynamically applied to bind multiple two-dimension hierarchies in line with multiple inheritances. Selective bindings to the default valuations across multiple hierarchies signify customized exceptional valuations. One feasible structure is to dynamically join constraint-based nodes on a three-dimension hierarchy both horizontally and vertically. Such a multi-hierarchy representation model can express the personalized exception but requires modestly extra maintenance. On the two-dimension user interface, since a two-dimension tabular form does not exhibit a logical sequence, we propose a three-dimension (3D) representation structure for expressing the logic of a decision tree or table that takes individual preference into account. Figure 7 shows the display of multiple hierarchies in a 3D representation structure to express the logic similar to a decision tree. On a real user interface, the shaded boxes could be laid behind the white boxes in the direction perpendicular to the computer screen and can be navigated in sequence if so desired.

As an example of implementation, dimension *A*, shown in Figure 7, can be used to guide a user in perceiving a better systemic assessment whereas dimension *B* can be used to effectively show the results of aspect-focused choices. Such a representation structure is capable of expressing a decision tree resulting from multiple inheritances. The above illustration simply depicts some typical designs we have experimented in combining visual and non-visual components together to infuse heuristic intelligence into online UIs. Various implementations of our model for online decision support should be explored to fit diverse online customers.




Figure 7. Using a two-dimension user interface to representing the logic similar to a decision tree.

CLARIFICATIONS AND CONCLUSIONS

Our approach is bound to have limitations. First, although it is reasonable to expect that the effective knowledge representation of an online investment environment would help individuals conduct a sound investment strategy that yields a fair market return, it is nevertheless not intended to help individuals achieve a superior return on investment. The reason for such a seemingly moderate ambition is that the knowledge to be represented in our model is merely derived from public and/or historical data. We argue that in an efficient market, as soon as some opportunities become obvious, many professional participants with equally efficient access to the market are probably already in an advantageous position over most individuals in grasping the opportunities. As the Internet nearly accessible to everyone, the price of nearly every security in the market has been efficiently adjusted minute to minute to reflect its fair value. Subsequently, a bet on an above-average return based on publicly available information is generally considered a random choice if our model is so attempted.

Secondly, a high return is always accompanied with a high risk. Therefore, it is not our objective to achieve a higher return than the market as a whole. Rather, the purpose of improving knowledge representation is to assist an individual in maintaining a healthy portfolio that fits the individual’s long-term financial goals. When the price of a good stock rises sharply for some time, its associated risk and thus volatility likely also increase. Consequently, pursuing the value appreciation as the only objective would inevitably deteriorate other qualitative attributes of the portfolio and possibly lead the portfolio to deviate away from the model portfolio initially established.

Thirdly, the discussed learning methods are not intended to respond to the specific events quickly. Also, depending on the underneath algorithms, these methods might regressively encompass the latest data but would then result in an unstable representation that might confuse novice investors. Therefore, our view is that an intelligent investment website may only be able to guide an individual to avoid common mistakes rather than to suggest optimal decisions as an expert system usually does.

The main contribution of our research is that amid the critical constraints of an online investment environment, our research has relied on the knowledge representation of available information over the Internet, rather than on the use of an inference engine or some offline facilitation that may just ruin the purpose of online trading, to deliver tailored decision support. The resultant model exhibits how the adapting aptitude of a knowledge representation framework can be constructed in order to accommodate the dynamic nature of the financial market and personal financial goals. Therefore, our research exemplifies an innovative, as well as practical, approach to providing online investors with personalized decision support without spoiling the intrinsic advantages of online trading.

REFERENCES

1. Barber, B. & Odean, T. (2001). Online investors: Do the slow die first? Paper presented at the Rodney L. White Conference on Household Financial Decision-Making, March 2001.
2. Dalkey, N.C. (1988). Models vs. inductive inference for dealing with probabilistic knowledge. Lemmer, J.F., & Kanal, L.N. (Eds.), *Uncertainty in Artificial Intelligence*, Amsterdam: North-Holland, 3-11.
3. Charles Schwab (2004). Schwab Fundamental Advice and Signature Advice, retrieved from <http://www.schwab.com> on August 24, 2004.
4. Formica, A. & Missikoff, M. (2004). Inheritance processing and conflicts in structural generalization hierarchies. *ACM Computing Surveys*, 36(3), 263–290.
5. Genesereth, M. & Nilsson, N. (1988). *Logical Foundations of Artificial Intelligence*. Palo Alto, CA: Morgan Kaufmann.
6. Kauffman, R., Subramani, M., & Wood, C. (2004). Analyzing information intermediaries in electronic brokerage. Working paper, Information and Decision Sciences Department, Carlson School of Management, University of Minnesota.
7. Kim, H. (2002). Predicting how ontologies for the semantic Web will evolve. *Communications of the ACM*, 45(2), 49–55.
8. Larkey, P. (2002). Models of decision making: their uses and limitations. Proceedings of the 2002 Computational Analysis of Social and Organizational Systems Conference, 1–10.
9. Mathews, I. (2004). Strategies of online brokerage firms to remain viable and competitive in the present day market environment. *ABAC Journal*, 24(2), 26–35.
10. Merrill Lynch (2004). Holding Company Depository Receipts, retrieved from <http://www.holdrs.com/holdrs/main/index.asp> on August 26, 2004.
11. Mertens, S., Marius, R., & Yuliadi, E. (2004). An intelligent dialogue for online rule based expert systems. Proceedings of the ACM Ninth International Conference on Intelligent User Interfaces, 208-282.
12. Minsky, M. (1975). A framework for representing knowledge. In P.H. Winston (Ed.), *The Psychology of Computer Vision* (pp. 211-217). New York: McGraw-Hill.
13. Rajgopal, S., Venkatachalam, M., & Kotha, S. (2000). Does the quality of online customer experience create a sustainable competitive advantage for e-commerce firms? Research Papers (Paper No. 1666). Stanford Graduate School of Business.
14. Reiter, R. (1980). A logic for default reasoning. *Artificial Intelligence*, 3, 81-132.
15. Reiter, R. (1985). On reasoning by default. In R. J. Brachman and H. J. Levesque (Eds.), *Readings in Knowledge Representation* (Paper No. 23). Los Altos, CA: Morgan Kaufmann.
16. Rich, E. & Knight, K. (1991). *Artificial Intelligence* (Second Edition, pp. 510-514). New York: McGraw-Hill.
17. Sharpe, W.F., Alexander, G.J., & Bailey, J.V. (1999), *Investments* (pp. 186-190). New Jersey: Prentice Hall.