Competition Dynamics in a Virtual Fund Management System

Casey Dougal cdougal@gmail.com
David Merriman rnjedi@email.byu.edu
Jeffrey Humpherys jeffh@math.byu.edu
Sean Warnick sean@cs.byu.edu

Information and Decision Algorithms Laboratories, IDeA Labs
Departments of Computer Science and Mathematics
Brigham Young University, UT
http://idealabs.byu.edu

***Abstract*— This research explores the development of a virtual fund management system to benchmark investment controllers as algorithmic decision processes on live market data. While similar paper-trading competitions exist, the Tour de Finance system is unique in its use of a particular class of dynamical systems as a dynamic rating mechanism to help ensure that the system is fair and rewards real intelligence over dumb luck. A link is provided for those wishing to participate in the Tour by designing controllers to interact with the system as autonomous agents, each managing its own virtual fund.**

I. Introduction

Some of the most complex and interesting examples of dynamical systems are financial markets. Perhaps no other system, whether natural or engineered, has received as much focused attention to understand superior methods for making data-driven decisions in the face of intrinsic complexity and uncertainty. For those interested in understanding the capabilities and limitations of decision dynamics that transform data into intelligent action, finance is an excellent arena to develop and test new methods.

There are many characteristics that make finance such a fertile testbed for decision and control technologies. Not only are the markets data-rich and highly instrumented, but their dynamic behavior can be extremely complex, exhibiting phenomena such as stampeding, swarming, chaos, bubbles, crashes etc. [1], [2]. Moreover, this behavior is inherently multi-agent [3] and can be driven by asymmetric irrational investor psychology such as fashion dynamics [4] or the fact that many investors appear to be risk-seeking toward choices involving sure losses while being risk-adverse towards choices involving sure gains [5]. Investor perception of risk and reward, along with the aggregation dynamics of the multi-agent system, contribute to the complexity of behaviors and perceived uncertainty of financial markets.

Another interesting aspect of finance as an application of decision and control technologies is the fact that there is comparatively very little overhead in the development of new financial machinery and systems. That is to say, while developing new sensors or actuators for engineered systems typically involves significant implementation work, the implementation of new financial instruments is primarily policy driven. Examples of such financial instruments include:

- Derivatives with payments linked to the S&P 500 stock index, the temperature at Kennedy Airport, and the number of bankruptcies among a group of selected companies [6],
- Derivatives with payments linked to the weather in a particular city [7], [8],
- Derivatives with payments linked to the 2008 Presidential election [9].

Note that exotic derivatives such as these are not offered as gambling mechanisms for entertainment. Rather, they are carefully designed financial instruments engineered to offset the perceived risk of other investment decisions. For example, ski resorts may invest in weather derivatives to hedge against weather that would adversely impact their business. This ability to efficiently redesign market touch points for the investor adds an interesting degree of freedom to the decision problem, and the ease with which new mechanisms can be deployed suggest that increasingly investment science is truly becoming *information* science.

The natural entrance to this exciting world of finance as an application arena for decision and control technologies is the stock market. Stock markets are comparatively simple, yet they retain the essential characteristic of all financial markets, namely, that good decisions hinge on one's ability to predict the behavior of time series data. In spite of their relative simplicity, however, stock markets still represent a level of complexity in choosing which assets to include in an investment portfolio, and in which proportions, that makes the decision problem non-trivial.

As a result, we have built a virtual fund management system to enable the benchmarking of various algorithmic decision processes used for portfolio management. The idea is that each algorithmic decision process is coded as an agent that interacts with a central system to manage a fictitious mutual fund. The system operates as a multi-agent system with many funds operating simultaneously in a competitive environment we call "Tour de Finance," after the famous cycling competition. Each fund is assigned a fictitious amount of money to manage, and the agent decides how to invest its assets. The central system updates each fund's book according to its market performance using live market data.

This paper reports the development of the "Tour de Finance" system and invites members of the controls com-

munity to participate in the Tour by designing investment controllers to benchmark on the system. Moreover, we report our work in designing a novel benchmarking system that uses competition dynamics as a gauge with memory to ensure that the system is fair and only rewards those agents that truly exhibit more effective processes for transforming data into decisions. In particular, we show how our system deals with the problem of rewarding lucky decision strategies using a carefully designed dynamic process.

II. Background

Competitions and tournaments are well studied devices for realizing particular dynamical systems. The dynamics of these events change the characteristics of winners and losers, thereby changing the efficacy of the competition as a selection mechanism. We are interested in using a virtual fund management competition as a mechanism to select algorithmic decision processes that are better able to predict dynamic behavior from data and use it for timely decision-making. Three related classes of such competitions include virtual investment systems, artificial markets, and trading competitions.

A. Virtual Investment: Paper Trading

Paper-trading systems have become extremely common with the advent of the internet and on-line trading systems. Many online brokers offer a paper-trading tool as a mechanism for potential investors to become familiar with the trading process and to explore the performance of different trading strategies. Nevertheless, these systems are primarily marketing tools, designed to bring traffic to a particular site, or customer service mechanisms, designed to help potential customers make the decision to buy a particular broker's services.

Paper-trading *competitions*, on the other hand, typically serve an educational purpose. Business schools, and even some high schools, have used such competitions to expose students to decision problems in finance and encourage them to explore the information necessary to support strong investment decisions [10], [11], [12].

A different example of a competition designed to understand investment dynamics is the Wall Street Journal Dartboard Contest. The idea was to empirically support or invalidate the long-contested efficient market hypothesis, which would suggest that no method of portfolio management is superior to another. From October, 1988, the Wall Street Journal has been comparing the performance of professional investors with that of portfolios managed through random throws at a dartboard covered in stock ticker symbols [13]. A new, independent competition was run each month for six months. In 1998 the Journal reported that the pro's won 61 of the first 100 contests against the darts, but only beat the Dow Jones Industrial Average (DJIA) 51 times. Although the pro's average gain was 10.8%, compared with that of 4.5% for darts and 6.8% for the DJIA, many criticisms have been levied against contest due to the lack of scientific rigor in the experiment design [14], [15]. One criticism suggested that the way the Journal measured performance was inadequate to draw conclusions about whether one fund was better than another. The lesson learned is that the design of a competition's dynamics is critical to one's ability to draw conclusions and use it as a research tool.

B. Automated Trading: Artificial Markets

Over the last decade there has also been an increasing interest within the finance community in describing equity markets through computational agent models. These studies focus on creating artificial markets in which automated trading agents are inserted and then left to act autonomously in an attempt to reproduce certain characteristics of their real-market counterparts.

An in-depth overview of earlier research in the agent-based computational finance field is presented in [16]. Overall, these models range from the very simplistic zero-intelligence random trader models of [17], to the highly detailed Santa Fe Artificial Stock Market described in [18]. The Santa Fe market exhibited agents with bounded rationality and inductive reasoning, and it was one of the first artificial markets to correctly simulate real world stock market dynamics such as bubbles, crashes, and continued high trading volume.

Essentially, these studies attempt to determine how real-market traders *actually* behave. The idea is to hypothesize a trader model and build it as an agent. The artificial market then exhibits certain emergent properties which, if matching real market behavior, might be considered as evidence that real traders behave similarly to the automated agents. Although this approach has been criticized [19], these examples demonstrate another way multi-agent competitions are being designed and used as an instrument to analyze market behavior.

C. Trading Competitions to Discover Good Strategies

While the competitions discussed so far focus on developing a model of the behavior of real markets and real investors, another class of competitions are designed to understand how investors *should* behave given the existing market dynamics. Two studies in this class are the popular Trading Agent Competition (TAC), conducted by the Swedish Institute of Computer Science (SICS) in cooperation with the e-Supply Chain Management Lab at Carnegie Mellon University [20], [21], and the Penn-Lehman Automated Trading (PLAT) Project, conducted by Micheal Kearns and Luis Ortiz of the University of Pennsylvania with additional support from the Proprietary Trading Group at Lehman Brothers [22].

TAC focuses mainly on multi-commodity auction simulations. For instance, in TAC-Classic, automated agents acting as travel agents compete against other travel agents in auctions for plane tickets, hotel reservations, and entertainment tickets in order to maximize the utility of the client they are working for while minimizing the cost of their client's trip. Instances of the competition run in 9 minute rounds, with the agents final score being the average over several instances of the competition. Overall, TAC deserves notice for the large

amount of industry and academic participation that it has received over its five years of existence.

The PLAT Project is built around the Penn Exchange Simulator (PXS), a software simulator for automated stock trading that merges automated client orders for shares with real-world, real-time order book data. So far, the project has focused on intra-day trading of one stock with each of the agents "cashing out" at the days end. Thus, the main focus of automated agents in the PLAT project is to use order book and other stock-related data to forecast the optimal time and volume to trade a single stock throughout a day to obtain a maximal profit. The PLAT project has had much success, also gaining participants from multiple universities.

Nevertheless, there are many aspects of the PLAT competition design that appear to be ad-hoc. For example, to prevent the results from being dominated by "lucky" strategies that simply place large bets in the form of excessive share positions, the PLAT project creates a rule that requires agents' share positions remain within a window of $\pm 100,000$ shares. Initially, in December 2003 and April 2004 performance in the PLAT project was measured solely by the Sharpe ratio [23]. Current PLAT competitions are scored by agents earning a fixed amount of points for achieving specific goals, such as daily profit and loss or daily intraday position reversals. This change in metric made a significant difference in the competition dynamics. Specifically, one agent in the competition, OBVol, despite having the highest Sharpe ratio and statistically significant profitability, finished fourth in its pool of 8 agents under the new metric. Although these changes seem reasonable, we couldn't find any theoretical justification for these rules and decisions in the literature.

III. Competition Design

To facilitate the Tour de Finance, we have constructed an online trading environment that allows participants to manage portfolios of stocks from the NYSE, NASDAQ, or AMEX exchanges. These portfolios can be managed either automatically, with pre-programmed trading strategies, or via human interaction. To regulate this competition, we have constructed a dynamical system which allocates market share to portfolios based on their ability to achieve consistently superior returns in the stock market. Thus, an agent's success metric in the Tour is a direct measurement of the agent's share of the fund-investor market, which is represented as a vector $z(k) = \begin{bmatrix} z_1(k) & z_2(k) & \dots & z_N(k) \end{bmatrix}$, where $z_i(k) \geq 0$ is the ith portfolio's percent share of the fund-investor market at time k, so that $\sum_{i=1}^{N} z_i(k) = 1$ for all $k = 0, 1, \dots$. By focusing on market share of portfolio investors, agents competing in the Tour de Finance are viewed as institutional fund managers, rather than individual investors, and market share becomes the ultimate metric distinguishing competitors.

A. Performance Measures

Measuring the performance of stock portfolios in a volatile stock market is a particularly difficult task. The basis for most investment performance measures begins with calculating the *one-period simple gross return*, which gives the percent return of holding an investment over one period. Defining $x_i(k)$, $i = 1, 2, \dots, N$, to be the total dollar value of the ith portfolio at time k, we can define the simple gross return of the ith portfolio at time k to be $\delta_i(k)$,

$$\delta_i(k) = \frac{x_i(k)}{x_i(k-1)}. \tag{1}$$

Traditionally, the one-period gross returns of a portfolio are measured over a period of time, and then they are used to compute performance measures such as average returns, or some risk-adjusted return such as the Sharpe ratio [24].

Average returns are calculated using either an arithmetic average $\frac{1}{k}\sum_{i=1}^{k} \delta_i(k)$, or a geometric average $(\prod_{i=1}^{k} \delta_i(k))^{\frac{1}{k}}$. Because the arithmetic average is greater than or equal to the geometric average, with equality occurring only in the instance that every $\delta_i(k)$ is the same for all k, it is a common practice of mutual fund managers to report the average annual rate of return as an arithmetic average. However, because both measures ignore the risk involved in the investment being measured they are inherently flawed as a tool for discriminating between intelligent strategies and merely lucky bets. Thus, a good return using these measures could indicate a solid investment, while it could also mean that a risky gamble paid off.

Better than these simple measures are risk-adjusted returns which attempt to measure the level of return on an investment relative to the amount of risk innate to the investment. Perhaps the most common risk-adjusted return is the Sharpe Ratio. The Sharpe Ratio defined as:

$$\frac{\bar{\delta}_p - \bar{\delta}_f}{\sigma_p} \tag{2}$$

where $\bar{\delta}_p$ is the average portfolio return, $\bar{\delta}_f$ is the average risk-free rate (usually proxied by the average return on short-term government T-bills or LIBOR), and σ_p is the standard deviation of the portfolio return series, all calculated over the same time period. Thus, the Sharpe ratio divides the average excess return of a portfolio by its standard deviation effectively measuring the reward to (total) volatility trade-off [24].

There are also many more risk-adjusted measures besides the Sharpe Ratio, including Jensen's measure [25], the Sortino ratio, Treynor's ratio, Modigliani-Modigliani measure, etc. However, because there exist so many different ways to define and measure risk, the various risk-adjusted measures which have been constructed often differ substantially and sometimes provide inconsistent assessments of performance [26].

Nevertheless, the problem with all of these risk-adjusted measures is that they require a model of the market to adjust for market risk. For example, the measures listed above are meaningful when the market follows the Capital Asset Pricing Model (CAPM). One reason for so many measures, then, hinges on the number of different models that exist for how people think the real market behaves, or what they think distinguishes intelligence from overly risky behavior. The real question, then, is whether it is *necessary* to employ

some model of market behavior to adjust returns in a way that distinguishes real intelligence from dumb luck.

We propose a measure that avoids using any model or belief of how stock prices change with time, or any assumptions about risk, etc. We do so by avoiding a discussion of risk altogether. Instead, we focus on the impact of a fund's performance on it's own investor market. Presumably, portfolio-investors will favor a smart fund over a lucky fund. Thus, from this perspective, *reputation* of the fund becomes the ultimate measure of a fund's performance. Reputation is a measure that adjusts for market share instead of risk. We define the reputation of fund i among N firms, $u_i(k)$, as the return of fund i weighted by its market share and normalized by the average return of all funds weighted by their market shares. That is, the performance of N funds at period k is the vector

$$u(k) = \begin{bmatrix} \frac{z_1(k)\delta_1(k)}{\sum_{i=1}^{N} z_i(k)\delta_i(k)} \\ \vdots \\ \frac{z_N(k)\delta_N(k)}{\sum_{i=1}^{N} z_i(k)\delta_i(k)} \end{bmatrix}, \tag{3}$$

with $\delta_i(k)$ and $z_i(k)$ the return and market share of fund i as previously defined. Thus, funds with equal market share have reputations at period k defined entirely by their relative returns. Nevertheless, as one firm grows market share over others, it gains a kind of inertia that requires larger shocks in its returns to change reputation, positively or negatively. This is a reasonable model of investor behavior since a dominating firm defines investor expectations, so no matter how it performs it sets the market average. Likewise, a firm with minimal market share also has a similar inertia since nobody is really paying any attention to how it performs. Using reputation as a performance measure, then, avoids modeling the behavior of the stock market and its associated risk by instead modeling the behavior of savvy portfolio investors that have to choose between N funds.

B. Competition Dynamics

While typical contests evaluate participants at the close of the competition, the Tour de Finance is dynamic, in that it is constantly regulating and evaluating the performance of its participants. To model the behavior of individuals in the portfolio-investor market, which regulates and scores the performance of competitors in the Tour de Finance, we build a model of how portfolio investors reallocate their investments between the N competing funds given new reputation information, $u(k)$. These competition dynamics are given by

$$Z(k+1) = AZ(k) + Bu(k) \tag{4}$$

where

$$Z(k) = \begin{bmatrix} z(k-n+1) \\ \vdots \\ z(k-1) \\ z(k) \end{bmatrix}$$

and

$$A = \begin{bmatrix} 0 & I & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & I \\ \alpha_1 A_1 & \alpha_2 A_2 & \dots & \alpha_n A_n \end{bmatrix} \quad B = \begin{bmatrix} 0 \\ \vdots \\ 0 \\ \beta B \end{bmatrix}$$

with $\alpha_1 + \alpha_2 + \ldots + \alpha_n + \beta = 1$ and $A_1, \ldots, A_n, B$ are row stochastic matrices. We will show in the next section that this dynamic system, driven by the reputation function, has certain properties that make it useful for governing the results of a competition regardless of its interpretation. Nevertheless, the competition dynamics (4) have a very intuitive interpretation modeling the dynamics of portfolio-investors choosing between N funds. Essentially, the model suggests that the market share vector at time $k+1$ is the convex combination of market share over the previous n periods with the current reputation vector $u(k)$. This allows the market share vector to evolve as a weighted moving average, with memory n. Typically we might think of the system as having fading memory, with $\alpha_1 \leq \alpha_2 \leq \ldots \leq \alpha_n$.

Moreover, the presence of matrices A_i, $i = 1, \ldots, n$ and B define the notion of rationality governing portfolio-investor behavior. For example, $A_1 = A_2 = \ldots = A_n = B = I$ would describe a situation where investors move their investment strictly based on performance. However, choosing a different set of matrices can model "advertising effects," or other irrational dynamics, that influence the relationship between reputation based performance, u, and the evolution of the market share distribution, Z.

The primary difference between the Tour de Finance and other competitions, then, centers on this reallocation dynamic. Instead of simply measuring performance at the end of the race, this dynamic system reallocates the book, or the total amount of money a fund has to invest, between funds based on their relative performance throughout the race. In doing so, we will show that the competition identifies consistently intelligent investment strategies and does not reward dumb luck, all without ever attempting to model the behavior of the stock market. In particular, our results do not depend on the CAPM or various definitions of risk.

Moreover, [27] empirically demonstrates that over the period of 1965-1984, mutual funds which recently and consistently performed well increased their share of the mutual-fund-investor market at the expense of poorer performing funds. This suggests that our competition model may be more than simply a dynamic rating mechanism for such competitions, but that it may be descriptive of the real-world effects that inspired its creation.

C. Portfolio Management as a Feedback Control Problem

By measuring fund performance from the perspective of the fund manager via our competition model in (4), the problem of managing a stock portfolio for each agent becomes a feedback control problem. Usually, the portfolio investment problem is purely open-loop: An investor uses time-series data to identify the dynamic system governing changes in stock prices, and then uses this information to

decide which stocks to buy and sell. The trades have no effect on states of the market, however, since the investor is typically small compared to the entire market capitalization.

Introducing reallocation dynamics, however, changes the situation considerably. This is because a fund may lose or gain money, not only by making poor investment decisions, but simply by under- or over-performing relative to competitors. Thus, although a fund does not have any specific information about its competitors, it can compute at each time step how its performance compares with the net competition based on whether money was reallocated above and beyond its market return. In this system, although a fund's investment decisions still do not affect the states of the market, they *do* strongly affect the states of the competition dynamics, Z. Moreover, it is these states that define winners and losers in the game.

To compete in Tour de Finance, then, an agent is designed as a feedback controller. At every time step, an agent has access to the amount of money it has to invest. This is the agent's only information beyond historical market data, such as daily return series (stock prices), yield, volume, etc. When an agent receives or loses funds because of relative performance, the money is added to or taken from a cash account and may be used in new investments. This account may go negative; there is no penalty for this, but negative cash amounts will prevent further investments. With this information, the agent needs to decide how to invest available funds to gain market share.

D. Characteristics of the Tour de Finance

There are a number of important properties about the dynamics of Tour de Finance that make it a provably good testbed for comparing control strategies. These include the following:

- The Tour de Finance is fair. There are a number of ways to discuss what it would mean for a competition to be fair. Here, we mean that if each of the agents
  the market dynamics are so unpredictable that virtually any investment strategy would perform, in some sense, equally well, then the system should converge to meaningful equilibria.
  - If the return of every fund is *identical* at each time step, $\delta_i(k) = \delta_j(k)$, $i \neq j$, and the n initial conditions are equal, $z(1-n) = \ldots = z(0)$, then $z(k) = z(0)$ for all $k = 1, 2, \ldots$. This follows immediately from the system (4) and has the intuitive appeal that, if all funds give identical returns, then investors have no incentive to move their money to another fund. Thus, whatever market share distribution initializes the system persists over time.
  - If, on the other hand, each fund's return is not necessarily identical, but behaves as a random sample from the same distribution i.e. has the same statistical properties, then the system converges to the uniform distribution for any non-trivial initial condition. That is, if the initial condition gives zero market share to all but one agent, then the market share will not evolve. Besides such pathological cases, however, things move toward uniform market share. Figure 1 shows this behavior from different initial conditions and with different return statistics.

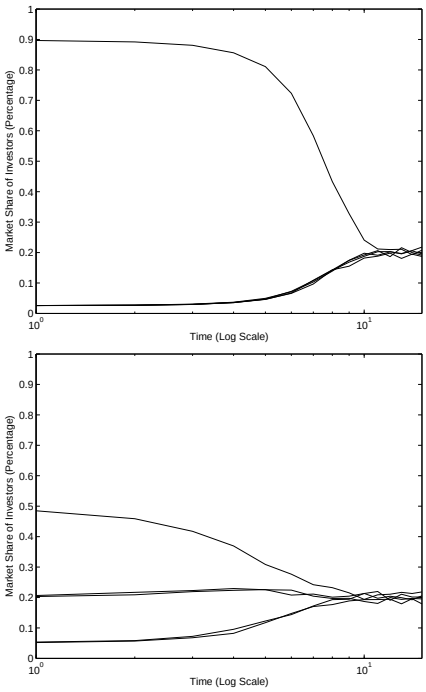

Fig. 1. Expected distribution of market share converges to uniform when returns between funds have the same statistics. These plots show the mean market share over 1000 trials for different initial conditions and return distributions. The top plot illustrates performance with log-normal returns, while the bottom plot simulates log-uniform returns.

    The fact that funds with identical return statistics converge to uniform market share, while funds with exactly identical returns maintain the initial market distribution, makes intuitive sense for a fair competition. The idea is that when the return statistics are identical, but the actual returns are not, investors will perceive different performance from day to day and move their investments accordingly. Because the performances are statistically equivalent, however, the expected distribution of investors will eventually equilibrate with equal market shares for all funds.
- The Tour de Finance rewards intelligence, not luck. Much of the motivation for using risk-adjusted performance measures is to avoid rewarding lucky decisions. Nevertheless, using models of market dynamics that may or may not be accurate to make the adjustment casts doubt on the meaningfulness of the conclusions. Tour de Finance accomplishes the same objective without hypothesizing what the dynamics of the actual stock market may be. Instead, consistently good strategies are favored by the reallocation dynamics over inconsistent ones. For example, in the simulation illustrated by Figure 2, Portfolio 1 has a lucky strategy which returns 40% in the first period but only 4%

thereafter. This is suggestive of a strategy that gets lucky and then invests in blue-chip stocks to protect against losses. The remaining portfolios are driven by lognormal disturbances. Portfolio 5 has a mean return of 10% and a variance of 17.7%, and Portfolios 2, 3, and 4 each have means of 6.7% and variances of 16.7%. Fixing the risk-free rate at 1% and calculating Sharpe Ratios, we see that Portfolio 1 has the best Sharpe ratio by a large margin. The Sharpe ratios are (in ascending portfolio order): 12.371 .128 .128 .128 .251. Nevertheless, the competition dynamics of the Tour de Finance system drive Portfolio 1 to the losing position because its lucky, inconsistent returns are soon forgotten. Instead, the steady average performance of Portfolio 5 dominates the market over time.

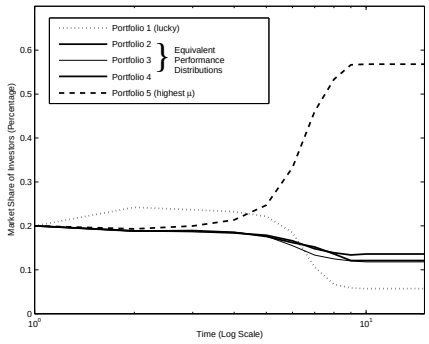

Fig. 2. Average market share for five portfolios over 1000 samples. While the Sharpe ratio rewards the lucky strategy of Portfolio 1, Tour de Finance punishes the inconsistent return and excessively timid strategy of Portfolio 1 in favor of the higher average return of Portfolio 5.

IV. CONCLUSIONS

We have presented a platform for detecting superior trading methodologies that consists of a stock-portfolio management competition regulated by an underlying dynamic system. Managing the competition in this way allows participants to attempt to control the system via their respective portfolio management strategies, thus transforming the optimal investment problem to a feedback control problem.

The characteristics of this system have been illustrated and its dynamics have been shown to be intuitive, fair, and useful. Those who are interested in participating in the competition are encouraged to contact us at sean@cs.byu.edu ATTN: Tour de Finance. The only prerequisite for participating in the competition is a computer with internet access. In particular, agents can be constructed using our C# or C++ API, or by using our pre-made Matlab shell. There is also an option to enter the contest by manually managing an agent by logging onto the server and executing transactions using the Tour de Finance's HTML forms interface. Additional information concerning the competition, how to enter and construct an agent, and when the next competition trial will begin can be found at http://tourdefinance.cs.byu.edu.

REFERENCES

[1] L. Veldkamp, "Media frenzies in markets for financial information," *The American Economic Review*, vol. 96, no. 3, pp. 577–601, June 2006.

[2] T. Lux, "Herd behavior, bubbles, and crashes," *The Economic Journal*, vol. 105, pp. 881–896, July 1995.

[3] B. LeBaron. (2006) Agent-based computational finance. Working Paper. [Online]. Available: http://people.brandeis.edu/ blebaron/wps/hbook.pdf

[4] E. McGoun, "Fashion and finance," *International Review of Financial Analysis*, vol. 5, no. 1, pp. 65–78, 1996.

[5] D. Kahneman and A. Tversky, "Prospect theory: An analysis of decision under risk," *Econometrica*, vol. 47, no. 2, pp. 263–292, March 1979.

[6] R. Stulz, "Demystifying financial derivatives," *The Milken Institute Review*, pp. 20–31, Third Quarter 2005.

[7] Y. Yamada, "Controlling business risks using weather derivatives," in *Proc. IEEE American Control Confernce (ACC'06)*, Minneapolis, Minnesota, USA, June 2006.

[8] L. Zeng, "Pricing weather derivatives," *Journal of Risk Finance*, pp. 72–78, Spring 2000.

[9] (2006, September) Iowa electronic markets. Henry D. Tippie College of Business, The University of Iowa. [Online]. Available: http://www.biz.uiowa.edu/iem/

[10] (2006) Stock market research analyst competition. UCLA, Anderson School of Management. [Online]. Available: http://www.anderson.ucla.edu/x2307.xml

[11] (2006) Redefining investment strategy education (rise) competition. University of Dayton. [Online]. Available: http://sba.udayton.edu/rise/index.asp

[12] (2006) Stock investment competition. University of Toledo. [Online]. Available: http://www.business.utoledo.edu/COBA/StockTradingCompetition/HSStockCompetitionInfo.asp

[13] J. Georgette, "Investment dartboard: Champion stock pickers ignore turmoil, select issues for 100th dartboard contest," *Wall Street Journal*, p. C.1, October 7, 1998.

[14] B. Liang. (1996, November) The "dartboard" column: The pros, the darts, and the market. Working Paper. [Online]. Available: http://ssrn.com/abstract=1068

[15] B. Barber and D. Loeffler, "The dartboard column: Second-hand information and price pressure," *Journal of Financial and Quantitative Analysis*, vol. 28, no. 2, pp. 273–284, June 1993.

[16] B. LeBaron, "Agent based computational finance: Suggested readings and early research," *Journal of Economic Dynamics and Control*, vol. 24, pp. 679–702, 2000.

[17] D. Gode and S. Sunder, "Allocative efficiency of markets with zero intelligence traders," *Journal of Political Economy*, vol. 101, pp. 119–137, February 1993.

[18] R. Palmer, W. Arthur, J. Holland, and B. LeBaron, "Artificial economic life: A simple model of a stock market," *Physica D*, vol. 75, pp. 264–274, 1994.

[19] D. A. Lane, "Artificial worlds and economics, part i," *Journal of Evolutionary Economics*, vol. 3, no. 2, pp. 89–107, May 1993.

[20] (2006) Trading agent competition. [Online]. Available: http://www.sics.se/tac/page.php?id=1

[21] M. Wellman, P. Wurman, K. O'Malley, R. Bangera, S. Lin, D. Reeves, and W. Walsh, "Designing the market game for a trading agent competition," *IEEE Internet Computing*, pp. 43–51, March/April 2001.

[22] M. Kearns and L. Ortiz, "The penn-lehman automated trading project," *IEEE Intelligent Systems*, pp. 22–31, November 2003.

[23] A. Sherstov and P. Stone, "Three automated stock-trading agents: A comparative study," in *AMEC*, 2004, pp. 173–187.

[24] W. Sharpe, "The sharpe ratio," *The Journal of Portfolio Management*, pp. 49–58, Fall 1994.

[25] M. Jensen, "Risk, the pricing of capital assets, and the evaluation of investment portfolios," *Journal of Business*, vol. 42, pp. 167–247, April 1969.

[26] Z. Bodie, A. Kane, and A. Marcus, *Investments*, 6th ed. New York, NY: McGraw-Hill/Irvin, 2005.

[27] R. Ippolito, "Consumer reaction to measures of poor quality: Evidence from the mutual fund industry," *Journal of Law and Economics*, vol. 35, no. 1, pp. 45–70, April 1992.