Towards Autonomous Investment Analysts - helping people to make good investment decisions

Paulo Andre Lima de Castro[1,2] and Ronnald Annoni Junior[2]

Autonomous Computational Systems Lab[1]
Aeronautics Institute of Technology (ITA)
São José dos Campos , SP, Brazil
Email: pauloac@ita.br

Raifa Intelligent Systems[2]
São José dos Campos, SP, Brazil
Email: annoni@raifa.com.br

Abstract

Since early days of computer science, researchers ask themselves where is the line that separates tasks machine can do from those only human beings can really accomplish. Several tasks were pointed as impossible to machines and later conquered by new advances in Artificial Intelligence. Nowadays, it seems we are not far from the day when driving cars will be included among the tasks machines can do in an efficient way. Certainly, even more complex activities will be dominated by machines in the future. In this paper, we argue that investment analysis, the process of assessment and selection of investments in terms of risk and return, should and can be among the tasks performed efficiently by machines in the (maybe not so far) future. Investment decisions have to be faced not only by financial professionals but by all people. Naturally, these professionals have more complex and often decisions to make, but everybody needs to invest to warrant good standard of living in the old age. In fact, there is significant research effort to create algorithms and/or quantitative methods to analyze investments. We present a brief review of them. Through this review, we may realize that there are many interconnected challenges in the quest for autonomous investment analysis. In this paper, we propose an adaptive multiagent architecture that deals with these three dimensions of complexity (nature of assets, multiple analysis algorithms per asset and horizon of investment) and keeps an explicit model of investor's preferences. This architecture breaks down the complexity faced by AIA in problems that can be addressed by a group of agents that work together to provide intelligent and customized investment advices for individuals. We believe that such architecture may contribute to development of AIA that deals with the complexity of the problem in a tractable way. Furthermore, this architecture allows the incorporation of known algorithms and techniques that may help to solve part of the issue.

Keywords—multiagent systems, automated analysis, technical analysis, financial analysis.

I. Introduction

In his great paper "Can a machine think?" [1], Turing discuss many objections pointed out to reinforce the idea that machines will never be able to really think. Some of these objections would not likely be raised against the idea of machines that can analyze investment, for instance the theological objection. It seems unlikely that someone would argue that analyzing investments is a function of man's immortal soul. However, it is interesting to observe that some objections could be raised against the possibility of autonomous investment analysts. We address here the more important and try to adapt Turing's arguments to show that these objections are not hard evidence of impossibility of an AIA in the future. Some other objections could be raised but we believe these are the most relevant. The chosen objections are explained and discussed below.

- The ***Heads in the Sand*** *objection*: The consequences of machines controlling investment would be too dreadful. The objection is more often formulated in a more subtle way: the risks of machine controlling investment would be too high. It would expose people to the possibility of lossing all their saving or even create catastrophic crises in global markets. It could also be pointed out that it could lead to unemployment of many people. In his paper, Turing despises such objection, stating that this argument is not sufficiently substantial to require refutation and consolation would be more appropriate. We would add that if AIA can be really efficient, perhaps it would be more likely that financial crises and bad investments would become less often. We discuss it in more detailed in section I-B.

- The ***Mathematical*** objection: Investment analysis or management is more than logic, it is kind of art, so it is beyond the limits of computability. It is well known in computer science that there are limitations to the powers of Turing machines (or simply limitations to what is computable). This objection to AIA is perhaps the closest to the original objection about machines that can think. We use here Turing's short answer and suggest the reader to refers to the original paper for a deeper discussion about this objection. Turing: "...although it is established that there are limitations to the powers of any particular machine, it has only been stated without any sort of proof, that no such limitations apply to the human intellect...Whenever one of these machines is asked the appropriate critical question, and gives a definite answer, we know that this answer must be wrong, and this gives us a certain feeling of superiority. Is this feeling illusory?...We too often give wrong answers to questions ourselves to be justified in being very pleased at such evidence of fallibility on the part of the machines. Further, our superiority can only be felt on such an occasion in relation to the one machine over which we have scored our petty triumph. There would be no question of triumphing simultaneously over all machines. In short, then, there might be men cleverer than any given machine, but then again there might be other machines cleverer again, and so on...".
- ***Arguments from Various disabilities***: This objection usually take the form: I believe you can make machines that do significant part of the job, but no machine will ever be able to do X. Numerous features X can be pointed out, for instance: be intuitive, have common sense, tell right from wrong, be innovative, think something really new. In fact, some of this features can be very hard to achieve, but the point is that no support is offered for these statements. We believe they are mostly founded on the principle of scientific induction. Nowadays, any person has seen many machines in her lifetime. From what she sees of them, she draws some general conclusions. They are useful for a very limited purpose, when required for a little different purpose they are useless and so on. From this observation, one may conclude that these are necessary properties of all machines. However, such conclusion is misguided by the assumption that new machines will have the same limits. In fact, Russell and Norvig state that during the early years of AI, AI researchers responded that claim by demonstrating one X after another [2] (p.17). It is true that some of the X given here are still to be demonstrated, as for instance: have common sense. In fact, common sense reasoning is one the branches of AI. However, there is no hard evidence it is impossible. Furthermore, one may argue that would be possible the existence of an efficient AIA even without common sense, provided it has access to all relevant information related to its target assets.

A. Why bother about autonomous investment analysis or management?

One really relevant question is why care about search towards autonomous investment analysts, if there are so many simpler tasks yet to be added to the list of things machines can do. We argue that investment analysis, the process of assessment and selection of investments in terms of risk and return, or even investment management (when it is delegated to the machine the power to execute the investments) should be among the tasks performed efficiently by machines. Some features make the problem specially interesting for autonomous approach: it is relevant to all people (**universality**), the presence of **conflict of interests** among analysts/managers and their investors is an open and complex issue and the high costs associated to investment management or analysis may make it unaffordable to many people. Let's discuss the two first issues a little more.

- **Universality**: Investment decisions have to be faced not only by financial professionals but by all people. Naturally, these professionals have more complex and often decisions to make, but everybody needs to invest to warrant good standard of living in the old age.
- **Conflict of interests**: The most common conflict of interest in a modern company happens among managers and stockholders (see [3], pg.12). However, there are possible conflicts of interest among analysts and investors, when analysts have investments on target assets themselves or are contracted by securities emitters. In fact, SEC (U.S. Securities and Exchange Commission) has a long history of examining potential conflicts of interests among such roles, for more information see [4]. For a more complete review of the theme, see [5]. Due to the fact that machine can have controlled or at least formally verifiable interests, possible conflict of interests can be avoided or at very least controlled in a more efficient way.

B. The future: What happens if autonomous investment analysts or managers become ubiquitous?

One interesting question about a future with AIA taking most of investment decisions is what would happen with average returns. Would it everybody in the world become rich or at least present very high average returns on their investments? The short answer is no. We believe the scenario described by Fama [6] in his Efficient Market Hypothesis (EMH) would take place. The EMH states that financial markets are efficient in pricing assets. Asset prices would reflect all information publicly available and the collective beliefs of all investors over the foreseeable future. Thus, it would not be possible to overcome the performance of the adjusted market risk, using information that is known to the market, except for simple chance. Thus, in the long-term, all investors would have returns limited to the performance of the market itself. Therefore, management of a portfolio would be simply the purchase of the market portfolio and keep it. Since, it would be impossible to consistently outperform the market. This line of action. It is usually called passive management, as opposed to the idea of managers trading assets in order to achieve higher performance than the market average. This alternative line of action is commonly called active management. The Efficient Market Hypothesis defines three different types of efficiency:

- **Market efficiency in weak form**: It implies that you cannot get return higher than the market only through the analysis of historical price series, although you can find assets that are priced overrated or under-valued in relation to the company's fundamentals that issued it.
- **Market efficiency in semi-strong form**: Asset prices adjust to their intrinsic values in short time and without overvaluation or overvaluation when new relevant information is made available to the definition of prices.
- **Market efficiency in strong form:** asset prices reflect all information available and adjust immediately to new information.

According to EMH supporters, the use of active management could be effective in weak form, however it would not be possible to obtain higher returns to the market consistently: such returns are possible only for short periods and chance. Therefore, they do not accept as examples rebuttal administrators assets that achieved superior return to the market, even in relatively long periods.

The controversy between EMH and active management supporters has produced many articles and books, including some dedicated to the general public, such as "A Random Walk Down Wall Street" [7] and "Beating the Street" [8] which support EMH and active management, respectively. Nonetheless, if AIA become ubiquitous it seems clear to us that markets would be much closer to EMH in its strong form, because asset prices would very fastly reflect all information available given the actions of such machines.

II. Related Work

The ultimate goal of an investment analyst, automated or not, is to find out and adopt the most desirable set of assets for an investor, according to his preferences. The analyst may adopt one set of assets through the *submission of buy and sell orders* to the stock market. The buy and sell transactions and price formation are defined through the processing of the orders of all investors in the market. In this section, we briefly present the two types of analyses used in asset management: technical and fundamentalist analyses (section II-A)

A. Technical and Fundamentalist Analyses

In the asset management domain, there are many analytic strategies based on time series analysis, which are often grouped in an approach called *technical analysis*. These strategies use some market information to identify patterns and to define orders. Some examples are *moving average*, *moving average converge-divergence*, *stochastic* and *relative strength index (RSI)*, but there are many others strategies [9].

Another approach to trading strategies is called *fundamentalist analysis*. It is based on information related to economic fundamentals (including company, sector and macroeconomic fundamentals), such as net profit, market share, revenues, sector growing rates, global growing rate among others. The fundamentalist analysis approach is less used in automated asset management, despite the fact it is widely used by human asset managers. Probably, this choice is due to the greater complexity to represent algorithmically many fundamentalist concepts.

Even within technical analysis, the identification of which information is really used and how the deliberation process occurs may change dramatically among different strategies. Furthermore, strategies may present very different performance according to market scenario [10]. This observation brought one first guideline to our architecture, i.e., to facilitate the composition of different strategies, like shown in section III.

Autonomous investment analysis is closely related (but not necessary equal) to many concepts, such as automated asset management, automated stock trading, algorithm trading, high frequency trading and some others terms. These terms have been a focus for many researchers [11], [12], [13], [14], [15], [16]. Among those papers, it is possible to identify two different groups, according to the typical time interval between orders (or position-holding period). Strategies that have to deal with short time intervals, like weeks, few days or even fractions of second cannot be based on fundamentalist analysis, because this latter is focused on long period scenario and it is can be used only when the typical time between orders are months or years. Therefore, technical analysis is widely used for short time intervals. When the holding-position period is very short, less than one day, it is often called high frequency trading. Many researchers and practitioners have been developing algorithms to achieve better performance exploring the fact that an automated system can analyze a significant larger amount of information when compared to a human being in small time periods [17], [18].

B. The environment faced by the autonomous investment analyst

Through this review, we may realize that many of these papers use historical series of price and/or volume to make inferences about investment decisions. This use of historical series to forecast future prices is controversial. Despite that, this practice, usually called technical analysis, is also widely used for analysts, at least as part of a more complex process of analysis that includes also market and economic information, such as profit, market share, EBITDA, price/profit ratio and so on. Methods that use company, market and/or economic information are commonly classified as fundamentalist. In fact, there are many papers that present trading algorithms based on **technical or fundamentalist information** and on some artificial-intelligence technique.

We may also realize that these algorithms have some parameters. However, there is not a good understanding of how the value of each parameters impacts the algorithm's performance and how the change of a parameter value impacts the setting of other parameters. This makes it very hard to define such values even for a small set of parameters. Furthermore, financial markets are **non-stationary environments**, i.e. the probabilities distribution may change along the time. Therefore, a specific algorithm may present great performance in a given period of time, but a terrible performance in the next period. Moreover, different assets may require different information and algorithms. For instance, if oil companies are very sensitive to changing of petrol prices, we cannot say the same about banks.

The environment faced by the autonomous investment analyst could be classified as: partially observable, sequential, stochastic, dynamic, continuous and multiagent, according to Russell and Norvig taxonomy [2], which is the most complex environment class pointed by them. However, it does not really represents the whole complexity of the problem. More than stochastic, such environment is also non-stationary process (the probability distribution do change along the time) and it is also strategic in the sense that two active investors compete for a more accurate valuation of assets and their acts may change other agents behavior.

Furthermore, a specific method of analysis may present great performance in a given period of time, but a terrible performance in the next period. Moreover, different assets may require different information and algorithms. For instance, if oil companies are very sensitive to changing of petrol prices, we cannot say the same about banks. We discuss how to measure analyst's performance in section III-C.

Besides the cited questions: definition of relevant information, non-stationary process and different nature of assets, we may also observe different requirements according to other dimension: the horizon of investment. We use this term to refer to the period of time the investor intends to keep his resources invested in the same set of assets. It may vary from several years to few milliseconds. This wide range leads to algorithms that can be very effective in very short horizons, but present poor performance in the long run. Another aspect that should be addressed by any Autonomous Investment Analyst (AIA) is that people do not have the same preferences about investments. Some investors may present a much stronger risk aversion than others, for example. An AIA must be aware of its investor's preferences, in order to provide appropriate advice.

The remaining of this paper is organized as follows: section II-A reviews some related work. In section III, we discuss an adaptive multiagent architecture that deals with this environment's complexity (different nature of assets, multiple analysis algorithms per asset and horizon of investment) and keeps an explicit model of investor's preferences. This architecture breaks down the complexity faced by AIA in problems that can be addressed by a group of agents that work together to provide intelligent and customized investment advices for individuals. We present a very simplified implementation of such architecture and analyze it in section IV. Finally, we propose some research directions to future work in section V.

III. COAST Architecture

The COAST architecture is designed to facilitate the simultaneous use of many analysis strategies and to explore the competition among these strategies, in order to achieve better results to the whole society. These strategies are materialized through agents called *advisors*. In COAST, strategies outputs are not interpreted as advices about one specific asset. There are three types of agents in COAST: **Advisors** which analyzes according to a strategy / specific formulation using a limited set of information about a specific asset; **Analyst** which gathers information from all the advisors of a particular asset and create an integrated analysis for a specific asset and **Manager** which is a specialist in the preferences of a particular

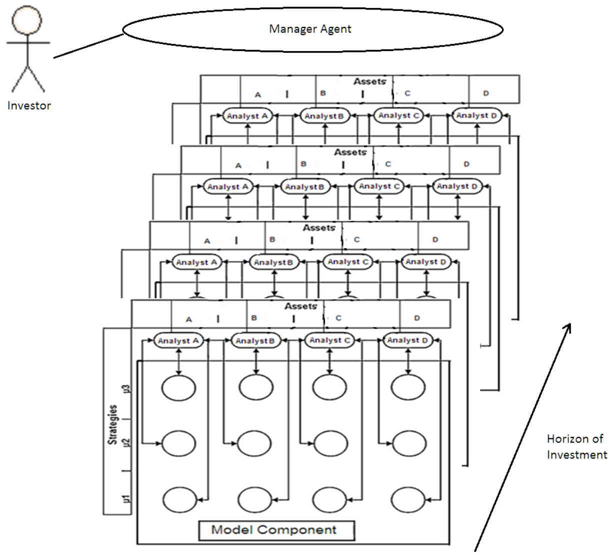

Fig. 1. COAST architecture extented to deal with several horizons of investment

investor and builds the best portfolio for that investor using the information of Analysts and its knowledge about the preferences of its investor.

The analysts and advisors agents are specific for one time horizon. For instance, there is an analyst for asset x and horizon one day and another analyst for the same asset in time horizon equal to one month. An investor operates with a specific horizon, although you can change it over time. There is one manager agent for each investor. The figure 1 shows an example of COAST architecture with four horizons of investment, four assets and three strategies per asset.

The other architectural guidelines are the following: **(i)** it should work with many different assets, **(ii)** it should adapt strategies' relevance for each asset and **(iii)** it should be able to deal with several different investor profiles. Each analyst is specialized in one single asset. Therefore, a society with four assets and three different strategies would be composed by four analysts and twelve advisors (three advisors for each analyst), like shown in the front plane of figure 1. The advisors located in a same column operate with the same asset and the analyst in top of the column evaluates and coordinates the work of the advisors in the column. For multiple horizons of investment, the architecture would have several planes of analysts and advisors. One plane for each horizon, figure 1 presents the case of four horizons of investment.

A. Advisors

Advisors suggest to buy or to sell a number of shares of a specific asset following their own strategy and their goal is to give the best possible advices to improve portfolio return. Advisors can be easily created using any well known trading strategy. These advices are sent to the analyst, who is the agent in charge of order definition. The advisor's life cycle is presented on the left of figure 2. In this figure, dashed lines show messages exchanged between agents and solid lines show state changes for each agent. Each state is represented by an ellipse, and has the following meaning :

1) **Asks for updated information**: The advisor, according to its strategy, asks for updated information from the *stock market simulator* [9], which can be seen in the center of figure 2.
2) **Receives information**: The stock market simulator returns the information which is locally stored. This step is also used to synchronize all the agents, in simulated time.
3) **Analyses and sends advice**: According to the collected information and his strategy, the advisor defines and sends a buy/sell/hold advice to its analyst.

The advisors performances are evaluated by their analyst according to their advices and the market evolution. For instance, whenever an advisor suggests buying an asset whose price arises after the advice, this advisor is positively evaluated. A similar reasoning can be made regarding a selling advice.

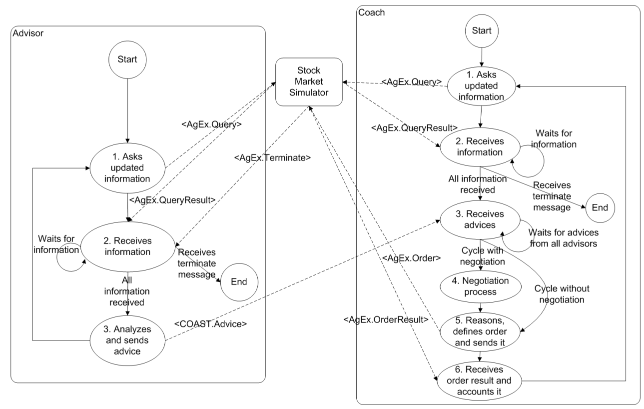

Fig. 2. Advisor and analyst life cycles in COAST architecture.

B. Analysts

Analysts basically *receive* advices, *evaluate* their advisors and *define* an aggregate estimate of asset price for its horizon, which is supplied to the manager agents. Analyst's decision process includes two stages: **filtering** eliminates some advisors (up to n-1, where n is the number of advisors) irrelevant in relation to the other and

R1. **If** *Advice* **is** *Buy* **and** *Evaluation* **is** *High*
Then *Expectation* **is** *Strong bullish*
R2. **If** *Advice* **is** *Sell* **and** *Evaluation* **is** *High*
Then *Expectation* **is** *Strong bearish*
R3. **If** *Advice* **is** *Buy* **and** *Evaluation* **is** *Medium*
Then *Expectation* **is** *Bullish*
R4. **If** *Advice* **is** *Sell* **and** *Evaluation* **is** *Medium*
Then *Expectation* **is** *Bearish*
R5. **If** *Advice* **is** *Hold* **Then** *Expectation* **is** *Unbiased*
R6. **If** *Evaluation* **is** *Low* **Then** *Expectation* **is** *Unbiased*

Fig. 3. Analyst expectation definition rules.

the **composition** phase which gathers information from various advisors to form a target price for the horizon.

1) Filtering of advisors: In filtering stage, the analyst receives data from all advisors and provides a list of 1 to N selected for the following composition stage, where N is equal or smaller then the number of advisors. The advisors elimination is done by discarding advisors with performance lower than an acceptable gap in relation to the best evaluated advisor. Such gap is a parameter for calibration (φ). The lower φ, the lower tends to be the selected group. If φ is 0, only the advisors with performance equal to the best one are selected. If φ equals to 1 the acceptable gap of performance is roughly 100

2) Composition of advices: Analysts calculate their estimate about price as a linguistic variable (expectation) with five terms: *strong bearish*, *bearish*, *unbiased*, *bullish* and *strong bullish* [19]. The **advices** that come from filtered advisors are also linguistic variables, with three linguistic terms: *sell*, *hold*, and *buy*. Finally, each advisor **evaluation** is also a linguistic variable and have three linguistic terms: *low*, *medium* and *high*, where the universe of discourse is the success rate of the advisor [0%-100%]. Each analyst receives advices from its advisors and follows the fuzzy rules presented in figure 3, in order to define its expectation.

The analyst expectation is used in **order definition**. The analyst defines his order according to his expectation using a fuzzy decoding method (in our implementation, center of gravity method). For instance, a *strong bullish* market leads to a buy order with high volume, meanwhile a *strong bearish* expectation leads to a sell order with high volume and a *unbiased* expectation makes the analyst to keep its current position.

C. Measurement of Performance

The purpose of the analyst of a specific asset is to make accurate price predictions. Thus, its performance should be measured by error rate or accuracy of its prediction. We use an *error rate* that has the following features: penalize more trend errors, belongs to the range 0-1, is independent on the horizon of investment. The *error rate* with such desired features is defined by the equation 1.

$$ErrorRate = \frac{1}{N} * \sum \left(\overline{P}_t - P_t\right)^2 + \alpha_{trend} \quad (1)$$

In equation 1, $\overline{P}_t$ is the estimated price for time t, P_t is the observed price for time t, α_{trend} is the weight of a trend error (estimate lower price, when it should be higher) and N is the number of data points in the given period of time. The same error rate can be used to measure performance of advisors of a given analyst.

D. Managers and Investor Profile

Each manager agent works with the horizon of investment frame given by its investor. It uses the analyses provided by the analysts specialized in that horizon. That is, a user with a horizon of one quarter changes for one month profile after two months thereafter, unless confirm that stays with three-month period. The manager agent filters and sorts the analyses according to the investor profile.

An user / investor has his profile defined by: time horizon, available capital, maximum acceptable risk and minimal acceptable return. The first two values may be directed informed by the investor. The last two are captured from him in a way that such information is kept within what is possible to obtain in the market. It can be done asking the user to choose from a range with a minimum in "very conservative" and maximum "very aggressive" (more natural) or a risk vs return chart.

IV. Simplified Implementation and Discussion

We have implemented a very simplified version of COAST architecture that uses four advisors strategies based on technical analysis and just one horizon of investment (one day). The technical indexes used are: moving average, moving average converge-divergence, relative strength index and price oscillator, mentioned in section II. We have selected 15 assets, which are part of the main index in Brazilian stock market, IBOVESPA, and presented a big number of trading in five years (from January, 2006 until December, 2010). We tested COAST societies trading in exchange using daily quotes, where each analyst could give one order a day. Additionally, we despised the effect of the orders given by the analysts in the market price, because the agents deal with a very small amount of money when compared to the traded volume for each asset. Despite the fact that the market simulator allows the use of transaction fees, for simplicity we set these fees to zero. In

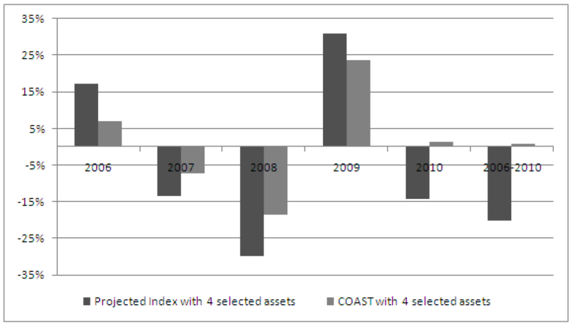

Fig. 4. Return achieved by COAST and *Projected Index* **using four selected assets**

fact, transaction fees have small influence on performance, since there is no big difference in the number of orders given by the analyzed societies [10].

In order to analyze COAST performance, we have executed simulation experiments using two different COAST societies, the first with four assets and the second one with all fifteen selected assets. Our first idea of direct comparison was to use the IBOVESPA index. However, a comparison among COAST performance and IBOVESPA would be biased because they do not deal with the same assets. In fact, IBOVESPA index composition changes in time and many assets have been included or excluded along the five years of the evaluation period, i.e. from 2006 to 2010. Due to these facts, we have created a theoretical portfolio called *Projected Index*, which is composed by the fifteen assets used by COAST societies, according to the relative weight of each asset (p_i). Moreover, we normalized these weights to use only the chosen assets using $p_i = \frac{w_i}{\sum_{j \in PI} w_j} * 100\%$, where p_i is the asset weight in Projected Index and w_i is the original weight in IBOVESPA. We have used these weights to define an AgEx trader agent [9], which buys and holds a set of shares according to the specified weights. This agent, called also *Projected Index*, acted in the same simulated evaluation period of five years. The figure 4 shows the performance of a COAST society managing four assets against Projected Index with the same four assets. In figure 5, we present the performance of COAST and Projected Index when managing all fifteen assets and finally the performance comparison among the two COAST societies are presented in figure 6.

The comparison among the COAST societies and the *Projected Index* agents in return (figure 4 and 5) shows a better performance of COAST in most of the years and in the whole period (2006-2010). In figure 6, we compare the performance of the COAST societies, the first dealing with four assets and the second with fifteen assets. It is easy to notice that the society with bigger number of assets presented a better performance in all years, except for 2006, and the best performance in the whole period.

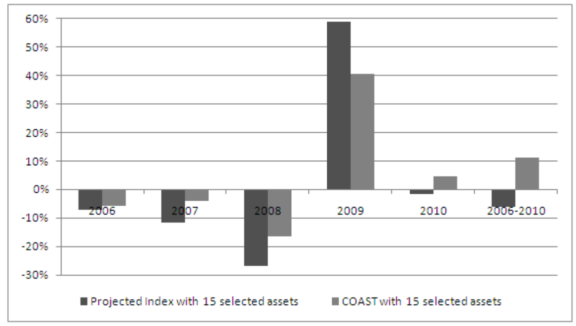

Fig. 5. Return achieved by COAST and *Projected Index* **using fifteen selected assets**

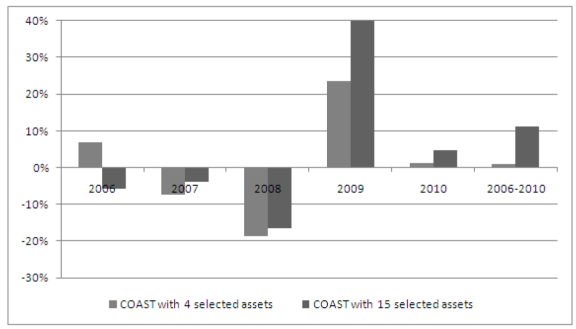

Fig. 6. Return achieved by two COAST societies with four and fifteen assets

These facts make us believe that it may be possible to pursuit better results with more assets and that is possible to achieve good performance in using a competitive agent approach.

V. Conclusions and Future Work

In this paper, we propose an adaptive multiagent architecture, which is populated by three types of agents: advisors, analysts and managers. Those agents deals with three dimensions of complexity (nature of assets, multiple analysis algorithms per asset and horizon of investment) and keeps an explicit model of investor's preferences. That way, this architecture breaks down the complexity faced by autonomous investment analysis in problems that can be addressed by a specialized types of agents. Those agents work together to provide intelligent and customized investment advices for individuals. We also presented a very simplified implementation of such architecture and tested in several simulation experiments, which results were presented and analyzed. In the simulated experiments, COAST architecture showed good results and overcome in some scenarios the chosen benchmark (*Projected Index*). It is also arguable that would be possible to achieve better results using more assets in society. We believe that such architecture may contribute to development of AIA that deals with the complexity of the problem in a tractable way. Furthermore, this architecture allows the incorporation of known algorithms and techniques that may help to solve part of the issue.

However, it is clear that there is a long path ahead to achieve an efficient autonomous investment analyst. A better understanding of how setting the several parameters in this multiagent architecture is needed. In fact as pointed out by LeBaron, a common criticism about Agent-based markets is that they usually have too many parameters and the impact of these parameters is not well understood ([20],pp. 1222). We also believe that a significant evolution would be a formal modeling of expectations, which are extremely important in financial/investment reasoning.

Acknowledgments

This work is partially funded by FAPESP. Grant 2013/50913-2, São Paulo Research Foundation (FAPESP).

References

[1] Alan M. Turing. Computing machinery and intelligence. *Mind*, 59(October):433–60, 1950.

[2] Stuart Russell and Peter Norvig. *Artificial Intelligence A Modern Approach Second Edition*. Prentice Hall, Englewood Cliffs - NJ, 2003.

[3] Aswath Damodaran. *Applied Corporate Finance*. Wiley, New York, NY, 2010.

[4] U.S. Securities and Exchange Commission. Analysts conflicts of interest: Taking steps to remove bias, 2016.

[5] Ingo Walter. Conflicts of interest and market discipline among financial services firms. In *Proc. of the Federal Reserve of Chicago conference on Market Discipline: Evidence Across Countries and Industries*, Chicago, USA, 11 2003. Fed-Chicago.

[6] Eugene Fama. Efficient capital markets:a review of theory and empirical work. *Journal of Finance*, 25:383–417, 1970.

[7] Burton G. Malkiel. *A Random Walk Down Wall Street*. W.W. Norton and Company, New York, 1975.

[8] Peter Lynch. *Beating the Street*. Simon and Schuster, New York, 1994.

[9] Paulo Andre Castro and Jaime S. Sichman. Agex: A financial market simulation tool for software agents. In Will Aalst, John Mylopoulos, Norman M. Sadeh, Michael J. Shaw, Clemens Szyperski, Joaquim Filipe, and Josï¿½ Cordeiro, editors, *LNBIP*, volume 24, pages 704–715. Springer, Berlin, 2009.

[10] Paulo André Lima Castro and Jaime Simão Sichman. Automated asset management based on partially cooperative agents for a world of risks. *Applied Intelligence*, 38:210–225, 2013.

[11] Keith Decker, Anandeep Pannu, Katia Sycara, and Mike Williamson. Designing behaviors for information agents. In W. Lewis Johnson and Barbara Hayes-Roth, editors, *Proceedings of the First International Conference on Autonomous Agents (Agents'97)*, pages 404–412, New York, 5–8, 1997. ACM Press.

[12] Kecheng Liu Yuan Luo and Darryl N. Davis. A multi-agent decision support system for stock trading. *IEEE Network*, 16(1):20–27, 2002.

[13] Michael Kearns and Luis Ortiz. The penn-lehman automated trading project. *IEEE Intelligent System*, 18(6):22–31, 11-12 2003.

[14] Alexander Sherstov and Peter Stone. Three automated stock-trading agents: A comparative study. In *Proceedings of the Agent Mediated Electronic Commerce (AMEC) Workshop - AAMAS 2004*, New York, 2004.

[15] Graham Kendall and Yan Su. Co-evolution of successful trading strategies in a simulated stock market. In *proceedings of The 2003 International Conference on Machine Learning and Applications (ICMLA'03)*, pages 200–206, Los Angeles, 6 2003.

[16] Xin Feng and Chang-Hyun Jo. Agent-based stock trading. In *proceedings of the ISCA CATA-2003*, Honolulu, USA, 3 2003.

[17] Irene Aldridge. *High-Frequency Trading: A Practical Guide to Algorithmic Strategies and Trading Systems*. Wiley, New York, 2009.

[18] Michael Durbin. *All About High-Frequency Trading*. McGraw-Hill, New York, 2010.

[19] Witold Pedrycz and Fernando Gomide. *An introduction to fuzzy sets: analysis and design*. The MIT Press, Cambridge, Massachusetts, 1998.

[20] Blake LeBaron. *Agent-based computational finance*. North-Holland, Amsterdam, The Netherlands, 2006.