# Who can I trust? Investigating Trust Between Users and Agents in A Multi-Agent Portfolio Management System

**Tiffany Y. Tang, Pinata Winoto and Xiaolin Niu**

Department of Computer Science
University of Saskatchewan
57 Campus Drive, Saskatoon, Saskatchewan, S7N 5A9 Canada
{yat751, piw410, xin978}@mail.usask.ca

**Abstract**

Trust between agents has been explored extensively in the literature. However, trust between agents and users has largely been left untouched. In this paper, we report our preliminary results of how reinforcement-learning agents (i.e. broker agents, or brokers) win the trust of their client in an artificial market I-TRUST. The goals of these broker agents are not only to maximize the total revenue subject to their clients' risk preference as most other agents do in [LeBaron et al. 1997; Parkes and Huberman 2001; Schroeder et al. 2000], but also to maximize the trust they receive from their clients. Trust is introduced into I-TRUST as a relationship between clients and their software broker agents in terms of the amount of money they are willing to give to these agents to invest on their behalf. To achieve this, broker agents must first elicit user models both explicitly through questionnaires and implicitly through three games. Then based on the initial user models, a broker agent will learn to invest and later update the model when necessary.

In addition to the broker agent's individual learning of how to maximize the 'reward' he may receive from his client, we have incorporated agents' cooperative reinforcement learning to adjust their portfolio selecting strategy, which is implemented in FIPA-OS. A large-scale experiment is expected as our future research.

## Introduction

Trust is a complex composition of many different attributes such as reliability, dependability, security, honesty, competence, timeliness, which may have to be addressed depending on the environment where trust is specified [Grandison and Sloman 2000]. It has been studied extensively in multi-agent systems [Jonker and Treur 1999; Elfoson 1998; Marsh 1992, 1994a, 1994b; Schillo et al. 2000], where trust is the "*attitude an agent has with respect to the dependability/capabilities of some other agent (maybe itself)*" [Jonker and Treur 1999]. However, we believe that trust is more than technical issues and should go beyond between agents, and also deals with trust between humans and agents. As Kini and Choobinch put it trust in a system involves "*a belief that is influenced by the individual's opinion about certain critical system features*" [Kini and Choobinch 1998]. This definition highlights human trust towards agents in electronic commerce, which motivates our study in this paper. In particular, we address the issue of the degree of trust a client has towards his broker agent to invest on his behalf: how can a client be sure that the broker agent will make a sound judgment based upon the risk-return preference of him. To describe our approach, let us first look at one real-life example.



**A real-life example**

Suppose Bank A offers customers a 24-hour on-line portfolio management service (investment in stocks, bonds, mutual funds, etc.) where one software agent will represent one customer to invest/manage his portfolio[1]. For Bank A, the overall goal is to attract as many customers as possible, and through that, receive higher revenues. Therefore, all the broker agents belonging to the bank are expected to cooperate to convince their respective clients that they are trustworthy[2]. As a result, customers will be willing to rely on them and 'dump' more money into the market. From the company's perspective, the cooperation between broker agents is the expected collective behaviors; whereas, for each individual broker agent, he must also do his own jobs to build trust between his client.

■

This real-life example highlights an important issue facing agent-based e-commerce systems: to win the trust of their clients and thus attract more customers to take advantage of the services they provide. It is the theme of our study in

1 A client can be an individual investor or a corporate investor representing a company.

2 Of course, for each broker agent, he might have his own self-interested goal to maximize his own revenue by building strong trust between him and his client(s), and through that, attracting more customers. But in the context of this paper, we only study the cooperation among agents to achieve a common goal.

this paper: investigating trust between software broker agents and their clients in the context of an artificial stock market called I-TRUST. Specifically, our study is based on a multi-agent portfolio management system—I-TRUST, where each broker agent represents a client to invest on his behalf in the market based on the client's characteristics (esp. risk-return preference). After each investment period, a broker agent will give an investment report to his client, and the client will evaluate and rate it. A broker agent needs not only to maximize the total revenue subject to his client's characteristics, but also has to learn the best portfolio selection strategy so as to attract his client to follow his expertise, e.g., invest more money in the next investment period. The higher amount of money a client is willing to put into the market, a higher degree of trust he has towards his broker agent.

Therefore, the degree of trust between a client and his broker agent is measured in terms of the amount of money a broker agent invest on his client's behalf. This trust can also be regarded as the degree of the client's reliability on the capabilities/competence of a broker agent and the service the agent controls or provides. It can be seen that this trust relationship is difficult to create; yet very easy to lose, therefore a broker agent has to learn to elicit his client's risk-return preference model so as to make a best investment strategy. From the perspective of the client, he implements a *partial* trust towards his client, which is different from most of current studies of this kind [LeBaron et al. 1997; Parkes and Huberman 2001; Schroeder et al. 2000; Decker et al. 1997]. In [LeBaron et al. 1997; Parkes and Huberman 2001; Decker et al. 1997], agents are *fully* delegated to make decisions in the artificial stock market, i.e. users hold *100% trust* towards the ability/competence of their agents; while in [Schroeder et al. 2000], users have *no trust at all* towards their agents: a broker agent acts only as an information agent to collect relevant information, and the 'intelligent client' will make final decisions on their own. Therefore, the client implements *no delegation* at all to his broker agents. However, we argue that it is reasonable to add another layer of delegation (*partial trust*) into the trust relationships, not only because it actually exists, but also because it can generalize trust into a wider and general spectrum by introducing trust between human and agents. In [Falcone and Castelfranchi 2001], the authors analyze trust, autonomy and delegation in multi-agent system and a framework is given for the theory of adjustable social autonomy in complex scenarios. Our reported work in this paper can be regarded as investigating trust, autonomy and delegation in a specific domain and expanding it when necessary.

The organization of this paper is as follows: in the next section, an overview of the system architecture is given. We will then describe the first part of our system: how to elicit user models as an underlying model for broker agents to act upon. And a detailed discussion of the reinforcement learning strategy a broker agent adopts is given. We will then describe our experiments and analyze the results obtained through a continuous simulation followed by some discussions. We conclude this paper by pointing out our future research directions.

### System Architecture

During the first step, users fill out a questionnaire and then play three games. The result of these activities results in an initial user model. Then, broker agents are thrown into the stock market and invest for their client through consulting with knowledge base (KB) storing synthesized stock information. The broker agents must adjust their portfolio-managing strategy based upon the feedback (including money delegated to the agent and instructions, if any) his client provides after each pre-defined investment period. Figure 1 illustrates the structure of I-TRUST.

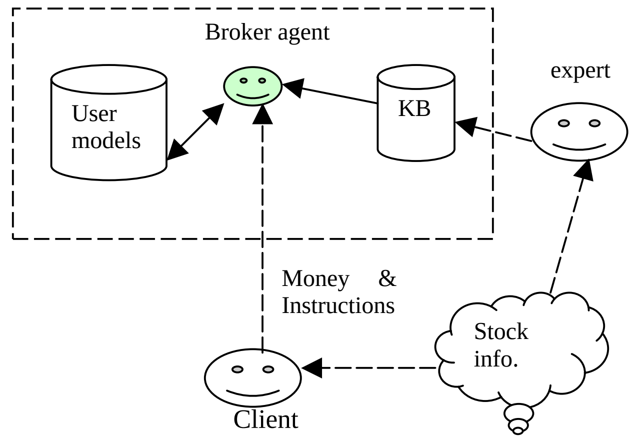


**Figure 1 System architecture-I-TRUST**

We assume that our artificial stock market only consists of few stocks and does not allow short selling or options.

## Eliciting User Models: Combining Questionnaires and Games

Building a client's model is crucial to both the decision making and decision supporting of a broker agent to act optimally; a model of a client's risk-return preferences is required. User models, in our domain, may seek to identify user's behavioral patterns and risk-return preferences. Generally it is assumed that user models are built manually, or at least the complete attributes necessary to build the user model are given as inputs. Thus, what the system needs to do is to calculate it based on some pre-defined predictive techniques. In other words, the expected utility of a client is *a priori* known (e.g., in terms of mean-variance) and this utility function is regarded as the user model based upon an optimal portfolio selection is made. However, the elicitation of user models needs a lot of information from the user. An alternative proposed in our system is to generate *initial* user models both explicitly through questionnaires and implicitly through game playing scenarios (currently we use three games only).

Then, a more detailed user model will be inferred automatically over multiple interactions between broker agent and his client. This is called collaborative-learning-based user modeling [Breese et al. 1998]. Generally, users are occasionally needed to fill in some forms giving out their preferences explicitly. However, this is not always accurate given user's misperceptions of their own behavior and the ambiguities inherent in any kind of qualitative elicitation of knowledge. Thus, in I-TRUST, besides the questionnaire, there are three game-playing scenarios related to our problem domain (i.e. stock portfolio management) for users to take part in. The game-playing scenario (with salient rewards) for eliciting users' preferences has been adopted extensively in the experimental economics community (cf. [Camerer 1995; Davis and Holt 1993]), and has proven to be a more accurate approach to infer users' preferences compared to a pure questionnaire [Roth 1995; Davis and Holt 1993]. The introduction of these game-playing scenarios to form an initial user model is a special feature of our system.

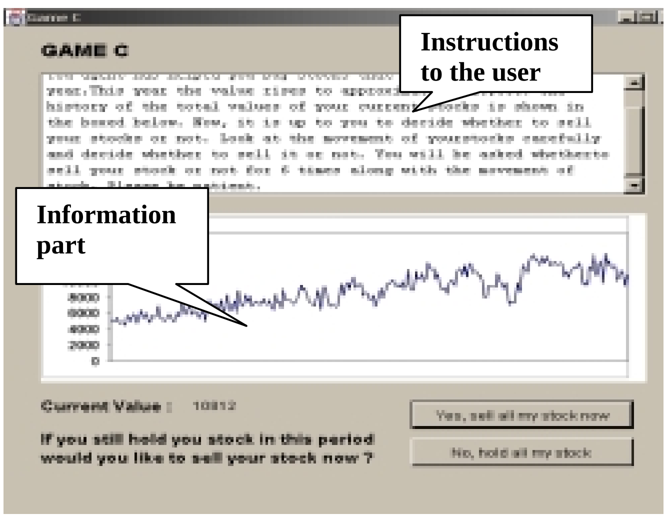


**Figure 2. A snapshot of Game C in our system**

Figure 2 shows one of the game-playing scenarios-game C. It is designed to investigate a client's *tolerance towards loss*. In particular, in time *t*, a client is asked to observe one stock's wave pattern (information part). He must decide whether or not he intends to sell his stock given this wave pattern. Six decisions should be made in this game. At one time, the stock plunged, but it is unclear that whether it is the lowest point or not. Thus, the user must decide whether he wants to sell at this time or not. (In a real-life scenario, a client is explicitly asked in questionnaires what is the maximum amount of money (in percentage) he will tolerate to lose). Formally, user model $U^j == \{R^j, EU^j\}$, where *j* represents the *j-th* client, combines the expected utility of this client (i.e. risk-return preferences) *and* the reward $R^j$ he gives out. Therefore, from the agent's perspective, he should not only maximize the total revenue and minimize risks for his client, but also maximize the reward collected, which in turn determines the portfolio selections and thus influences his client's expected utility $EU^j$. In our user model, the client's expected utility do not necessarily mean the von Neumann-Morgenstern expected utility but captures a broader notion of the mapping of his risk-return preference to portfolio preference. Agent elicits his client's risk-return preference from the initial questionnaire and games and then uses it to select portfolio. Ideally, we could categorize each user based upon his attributes, such as risk-return preferences, patience toward future income (inter-temporal substitution), tolerance toward loss, age group, income group, etc. In our current work, we combine all information into a one-dimensional metric: `user_type`. The formula is as follows,

$$\texttt{user\_type} = \mathbf{w'} \bullet \mathbf{x}$$

where **w** is the weight vector satisfying $\Sigma_i w = 1$ and **x** is the result vector, and $\texttt{user\_type} \in [0,1]$.

## Winning the Trust of Client: Single Agent Reinforcement Learning

In this section we will show how agents adopt reinforcement learning technique to learn to win the trust of his client. Specifically, how a broker agent analyzes the feedback from his client, and adjusts his portfolio selection strategy to collect higher rewards (trust) from the client. In this environment, the broker agent will learn what to do, i.e. how to map situations to actions, so as to maximize a reward from the environment (i.e. ratings from his client in terms of the amount of money the client is willing to invest in the next investment period) [Sutton and Barto 1998][3]. Figure 3 illustrates the single-agent-reinforcement learning process.

Note here that the dashed arrows show agents' interactions with their environment. Although a client might give inconsistent rewards to his agent over some investment periods, which might indicate that the client's risk attitudes might have changed. However, we will not consider it till it exceeds some threshold. It is consistent with real-life experiences: if a risk-averse-client becomes angry when he suffers a loss because he thinks that his agent is too cautious, it does not necessarily mean that he is becoming risk seeking. In I-TRUST, only after inconsistent rewards are significantly observed will the agent consider updating the user model.

---

[3] In MAS, reinforcement learning techniques have been studied extensively [e.g. Tan 1993; Stone 2000; Berenji et al. 2000; Stone and Veloso 2000; Arai et al. 2000].

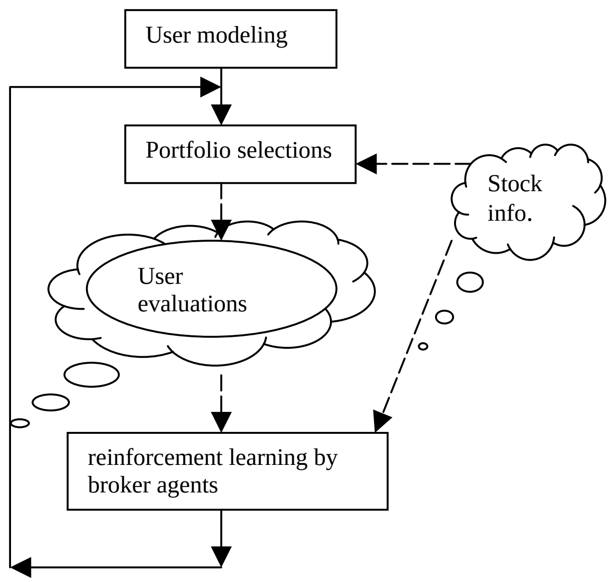


**Figure 3. A single-agent-reinforcement learning process**

## A Formal Model of A-Single-Agent-Reinforcement Learning

Let $\alpha \in [0,1]$ be a model parameter, and $P^a$, $P^u$, and $P^f$ represent the portfolio selection by agent, the client and the final adopted portfolio respectively. When $P^a \neq P^u$, agents will adopt convex combination of $P^a$ and $P^u$ to obtain $P^f$. The *convex combination policy* is

$$P^f = \alpha P^u + (1-\alpha) P^a$$

where if $\alpha \rightarrow 0$ then $P^f$ is an agent-dominated policy ; if $\alpha \rightarrow 1$, then $P^f$ is a client-dominated policy. $\alpha$ will be updated periodically by agent based on the reward gained from the client (reinforcement learning).

Let $r_p$ and $r_a$ denote the expertise reward and an investment reward respectively. Specifically, expertise reward $r_p$ represents the degree of trust a client holds towards the expertise of his agent, whereas $r_a$ is the reward a client is willing to give out in terms of the amount of money delegated to the agent. $r_p$ is subjectively updated based on the distance between $P^a$ and $P^u$ (an example will be provided in the next section).

The total reward obtained by an agent is a function of $r_p$ and $r_a$, i.e.

$$r = f(r_p, r_a)$$

In our current study, we choose $f(r_p, r_a) = (r_p + r_a)/2$.

We introduce the notion of investment period $t$, then we have:

$$r_t = (r_p + r_a)/2 \quad \text{and} \quad P_t^f = \alpha_t P_t^u + (1-\alpha_t)\, P_t^a$$

$$\alpha_t = \beta f(\alpha_{t-1}, r_t) = \beta(r_t - \alpha_{t-1}) + \alpha_{t-1}$$

In our current study, we choose learning rate $\beta = 0.5$, thus, $\alpha_t = (\alpha_{t-1} + r_t)/2$.

## Experimentation and Results

In order to verify our model, we start with three experiments for the purpose of the 'proof of concepts' (trust between humans and agents). Specifically, we shall attempt to investigate whether trust is build upon based on any of the following features or both:

- Client's trust in agent's expertise
- Whether the agent's portfolio selection strategy $P^f$ should be more client-oriented (i.e. based more on $P^u$ )

In order to investigate it, we design three experiments.

### Experiment 1 (Trust without Delegation)

In the first experiment, the final portfolio is selected according to the client, although his broker agent will also propose his own, which, of course, could be adopted by the client. After each round, a brief report is given out to the client listing the final profit made $\pi_t^f$, and the profit made by following agent's advice $\pi_t^a$.

Here, the measure of trust is based on the deviation of client's decision from agent's advice. For instance, if agent's opinion is to buy $5000 worth of stocks 'A' and its client only buy $2000 of it, then the trust level of his client is set to 0.7 (complete trust = 1, distrust = 0, it is determined subjectively). However, if the client chooses to sell all of his stocks 'A', then the trust level toward his agent is set to 0. Figure 4 shows a weak positive correlation between the profit made by following agent's previous advice, $\pi_{t-1}^a$, and the trust toward agent's next advice. Vertical axis represents the level of trust [0, 1] and normalized profit, and the horizontal axis represents investment periods. A strong correlation (= 1) means that client will absolutely follow agent's current advice if its previous advice made a high profit, or disregard current advice if its previous advice does not made profit (i.e. the level of profit influences level of trust). However, if the correlation closes to zero, then client's decision is not correlated to the fluctuation of profits. Our preliminary test with small sample size (n = 6) and *without* salient rewards shows an average correlation equals to 0.235 and an average trust level (in scale [0, 1]) equals to 0.72. These results should not be interpreted as statistically significant because of the small size of sample and the existence of bias due to an absence of salient rewards (surely people will think more seriously if they are playing with real money). However, it shows a reliance of clients towards their agents when they have relatively few knowledge of the stock portfolio. During our conversations with subjects after playing with the game, some people admit that they do not have expertise in portfolio management, and others who have some prior knowledge found that their agents are ‘smart’ enough, i.e. superior in making decision compared to them.

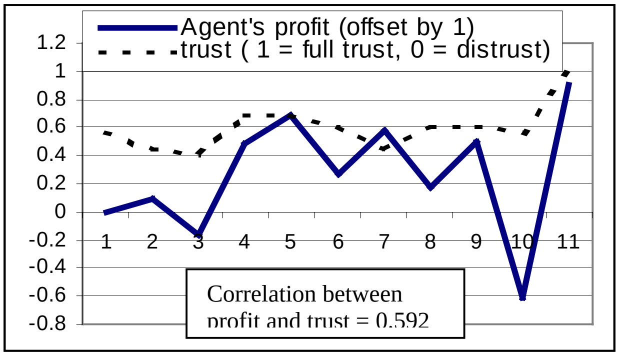


**Figure 4. One result of the first experiment**

## Experiment 2 (Trust with Full Autonomy)

In the second experiment, an agent selects the final portfolio and a client cannot intervene with agent’s decision. Agent’s decision is merely to find optimal portfolio investment based on client’s risk-return preference (user model) and market situation (environment). After each round, a brief report is given out to the client listing the final profit made (feedback). Then, the client decides whether to add/withdraw money to/from his agent or takes no actions. This is the only control implemented by a client. The measure of trust is based on how much money the client is willing to give to his agent in run-time adjustment. The result of our preliminary experiment is shown in figure 5 and 6.

It is shown from figure 5 that four out of six clients start to distrust their agents in period eight (when agents made small profit). And eventually two of them (User1 and User6) end up with completely distrusting their agents after period ten (when they are suffered some loss). Figure 6 shows two data sets representing the relation of trust and profit. From the figure we can see that User4 is persistent in his trust regardless of the profit made by his agent, while User6 withdraws all his money after his agent made negative profit (loss).

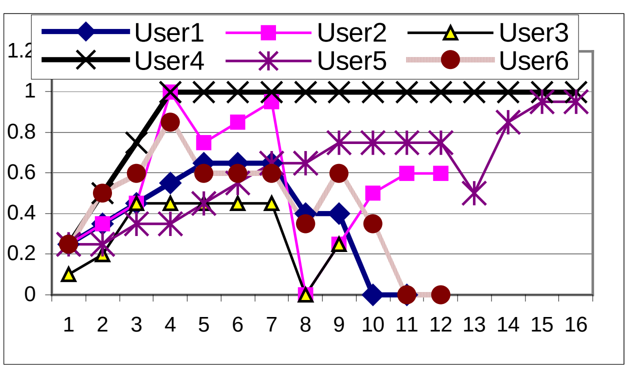


**Figure 5. Results of the second experiment**

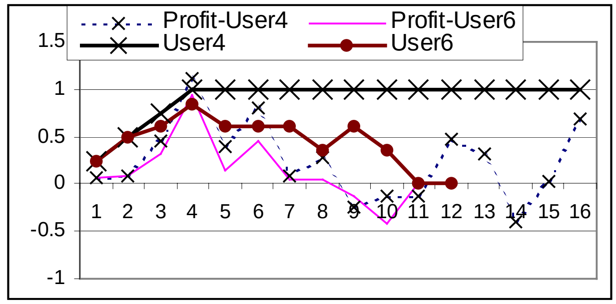


**Figure 6. Relation of the profit and trust observed in the second experiment**

From the result above we can infer that User6 is more risk averse compared to User4. This is true when we compare their user model elicited from questionnaire and games played before the experiment. User4 is a young man, patient in waiting future income, slightly risk averse, and has high tolerance towards loss. While User6 is an old man, less patience in waiting future income, slightly risk averse, but has no tolerance toward loss! The conformity of user model with the result convinced us that by retrieving user model, we could improve agent’s behavior to win more trust from the clients.

## Experiment 3 (Delegation with Mixed Control)

In the third experiment, a client delegates his agent a certain amount of money and then controls his agent’s behavior by adjusting the amount of money delegated and suggesting his preferred portfolio. However, an agent may not follow the client’s suggestion on portfolio selection but must follow his instruction on amount of money for the next investment period. Both of these would be interpreted by agent as the trust given by its client. Thus, agent will use the rewards to adjust his investment policy (in terms of $\alpha$ in convex combination of $P^a$ and $P^u$) by means of reinforcement learning technique described in previous section. Figure 7 shows one of the adjustment processes of $\alpha$ over time. The value of $\alpha$ is directly influenced by client’s rewards (feedback), therefore can be used to measure the trust. $\alpha$=0 is reached *iff* the client always agrees with his agent’s opinion *and* ‘dumps’ all his money to agent (trust = 1). And $\alpha$=1 *iff* the client does not agree with his agent’s opinion at all or withdraws all his money (trust = 0). From the third experiment the average value of $\alpha \approx 0.2$.

Unfortunately, most of the experiments only run for 10 until 20 periods, which is too short to contribute to the justification of the stability of learning parameters (long term relation.). However, we are convinced that this setting is better than settings in experiment 1 and 2 for two reasons:

1. Agent can learn client’s behavior (from client’s feedback), thus improve the flexibility of the system to prevent the breakdown of the client-agent relationship.

From our preliminary experiment, none of the users withdraw all his money as what happen in experiment 2. In practice, users in this experiment have more control compared to those in experiment 2, which in turn reduces their reluctance to use the system.

2. The autonomy of agent in decision making could help clients who are too busy (unwilling to intervene agent frequently) to manage their portfolio. And the system could reduce the risk of investment by considering both agent and client's opinions (using convex combination method), which in turn could reduce the error which might be made by either party (see figure 8).

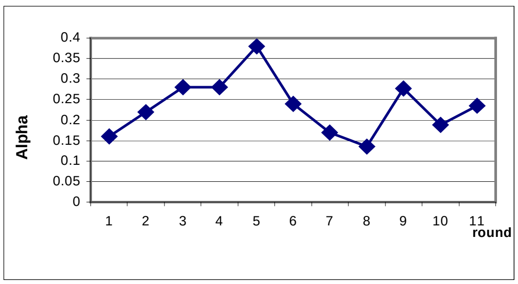


**Figure 7. The learning curve of agent in the third experiment**

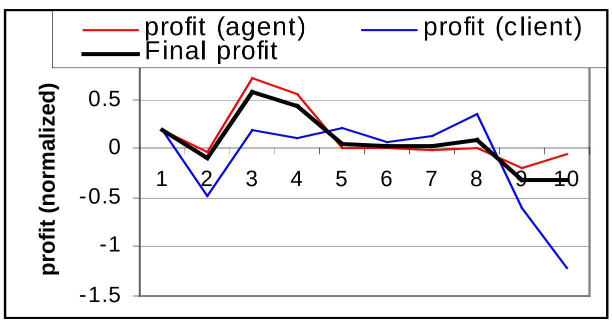


**Figure 8. Risk sharing observed in the third experiment**

## Discussions

• How broker agents cooperate to win the trust of their respective client.

In particular, what kind of information will be helpful for them to share? The portfolio episodic experience? Or just the portfolio selection strategy (it can be an instantaneous information or long-term learned knowledge)? There exist successful stories of the advantages of mutiagent cooperative reinforcement learning over independent learning [Tan 1993, Berenji et al. 2000]. However, the environment within which our study is conducted is different from the environments in [Tan 1993; Berenji et al. 2000]. In particular, more uncertainties have been observed, e.g., clients' can change their attitudes towards the tradeoff between profits earned and risks. For example, a client who is *risk averse* might be angry when he observes that his agent is too *cautious*, thus making him lose money, and the trust that has been built up might be destroyed instantly although the loss suffered can be gained in the future (because the client is not patient and switches to another agent)! Or a similar client who is *risk averse* could be very *happy*, because his broker agent makes a lot of money for him even if his agent invests in a rather *risky* asset, which in turn transfers a signal to the agent that the client might be willing to gamble occasionally! These attributes (i.e. clients' inconsistency attitudes towards profits and risks etc.) can be very *dynamic* (*concept drift* in machine learning [Widmer and Kubat 1996]). Therefore, it is essential to incorporate flexible cooperative learning algorithms capable not only of *capturing* but also of *adjusting* to these changes.

• To allow more than one broker agent to represent one client, so as to investigate competition as well as cooperation between agents in this domain

In this case, each broker agent will take chances to cooperate as well as compete to maximize the reward (trust) he collects, and he should also be smart enough to *anticipate when* to compete and *when* to cooperate.

## Future Works

As noticed from this paper, we only describe single agent's learning in I-TRUST. We have incorporated agents' cooperative learning to adjust their portfolio selecting strategy, which is implemented in FIPA-OS. A large-scale experiment is expected as our future works. Besides, all subjects *voluntarily* took part in the experiments, thus a bias is expected to be observed in the results of all experiments (without salient rewards). Therefore, we hope to conduct a large-scale experiment with salient rewards in our next study. However, even so, it is still too costly to have enough subjects with well-distributed risk-preferences attitude (i.e. risk averse, risk neutral, risk seeking) for the experiments. To overcome it, in our future study, we are considering the possibility of simulating a diverse range of artificial clients to ensure full coverage of the various behavioral patterns, and to augment the human subject studies.


### Acknowledgements

We are indebted to our supervisor Gordon McCalla for his many insightful comments, suggestions and encouragement; Pinata and Xiaolin's supervisor Julita Vassileva for her support, and valuable inputs. And we appreciate all the volunteers for taking their time to take part in our experiments. We would also like to thank the Canadian Natural Sciences and Engineering Research Council (NSERC) for funding Tiffany Y. Tang and Pinata Winoto through research assistantships and the University of Saskatchewan for funding Xiaolin Niu through a graduate teaching fellowship.

## References


Grandison, T and Sloman, M. 2000. A Survey of Trust in Internet Applications. *IEEE Communications Surveys*, Fourth Quarter.

Jonker, C.M. and Treur, J. 1999. Formal Analysis of Models for the Dynamics of Trust based on Experiences. In *Proc. of MAAMAW'99*, LNAI 1647, Springer-Verlag.

Arai, S., Sycara, K. AND Payne, T. R. 2000. Multi-agent Reinforcement Learning for Scheduling Multiple-goals. In *Proc. of the 4th International Conference on Multi-Agent Systems (ICMAS'2000).*

Berenji, H.R. and Vengerov, D. 2000. Advantages of Cooperation between Reinforcement Learning Agents in Difficult Stochastic Environment. In *Proc. of the 9th IEEE International Conference on Fuzzy Systems (FUZZ-IEEE'00*).

Breese, J.S., Heckerman, D. AND Kadie, C. 1998. Empirical Analysis of Predictive Algorithms for Collaborative Filtering. In *Proc. of the 14th Conference on Uncertainty in Artificial Intelligence, (UAI98)*, Madison, Wisconsin, 43-52.

Camerer, C.1995. Individual Decision Making. In *The Handbook of Experimental Economics*, Kagel, J.H. and Roth, A.E. (eds.), 587-703, Princeton University Press.

Davis, D.D. and Holt, C.A. 1993. *Experimental Economics*. Princeton University Press.

Decker, K., Sycara, K., Pannu, A. and Williamson, M. 1997. Designing Behaviors for Information Agents. In *AGENTS'97.*

Elfoson, G. 1998. Developing Trust with Intelligent Agents: an Exploratory Study. In *Proc. of the 1st Int'l Workshops on Trust*, 125-138.

Falcone, R. and Castelfranchi, C. 2001. The Human in the Loop of a Delegated Agent: the Theory of Adjustable social autonomy. In *IEEE Trans. on Systems, Man and Cybernetics, Part A: Systems and Humans*, 31(5), 406-418.

Kini, A. and Choobinch, J. 1998. Trust in Electronic Commerce: Definitions and Theoretical Considerations. In *Proc. of the 31st Hawaii Int'l Conf. System Sciences*, Hawaii.

LeBaron, B., Arthur, W.B., Holland, J.H., Palmer, R., and Tayler, P. 1997. Asset Pricing under Endogenous Expectations in an Artificial Stock Market. *The Economy as an Evolving Complex System II*, Arthur, W.B. et al (eds.), Reading, MA: Addison-Wesley.

Marsh, S. 1992. Trust and Reliance in Multi-agent Systems: a Preliminary Report. In *Proc. of the 4th European Workshop on Modeling Autonomous Agents in a Multi-Agent World (MAAMAW'92),* Rome.

Marsh, S .1994a. Formalizing Trust as a Computational ConcepT. PhD Thesis, Department of Mathematics and Computer Science, University of Stirling, UK.

Marsh, S 1994b. Trust in Distributed Artificial Intelligence. *Artificial Social Systems*, Castelfranchi, C. and Werner E. (eds), LNAI 830, Springer-Verlag, 94-112.

Parkes, D.C. and Huberman, B.A. 2001. Multi-agent Cooperative Search for Portfolio Selection. *Games & Economic Behavior*, Special issue on artificial intelligence and economics, 35:124-165.

Roth, A.E. 1995. Introduction to Experimental Economics. In *The Handbook of Experimental Economics*, Kagel, J.H. and Roth, A.E. (eds.), 3-109, Princeton University Press.

Schillo, M., Funk, P. and Rovatsos, M. 2000. Using Trust for Detecting Deceitful Agents Artificial Socities. In *Applied Artificial Intelligence* Journal, Special Issue edited by Castelfranci, C., Tan, Y., Fancone, R. and Firozabadi, B. on Deception, Fraud and Trust in Agent Societies.

Schroeder, M., McCann, J. and Haynes, D. 2000. Rough Traders and Intelligent Clients. In *Proc. of the International Conferences on Autonomous Agents (Agents'2000*), Barcelona, Spain, 259-260.

Stone, P. and Veloso, M. 2000. Multiagent Systems: a Survey from a Machine Learning Perspective. *Autonomous Robots*, 8 (3).

Sutton, R.S. and Barto, A.G. 1998. *Reinforcement Learning: an Introduction*. MIT Press.

Tan, M. 1993. Multi-agent Reinforcement Learning: Independent vs. Cooperative Agents .In *Proc. of the 10th International Conference on Machine Learning (ICML' 93)*, 330-337.

Widmer, G. and Kubat, M. 1996. Learning in the Presence of Concept Drift and Hidden Context. *Machine Learning* 23, 69-10.

# Appendix

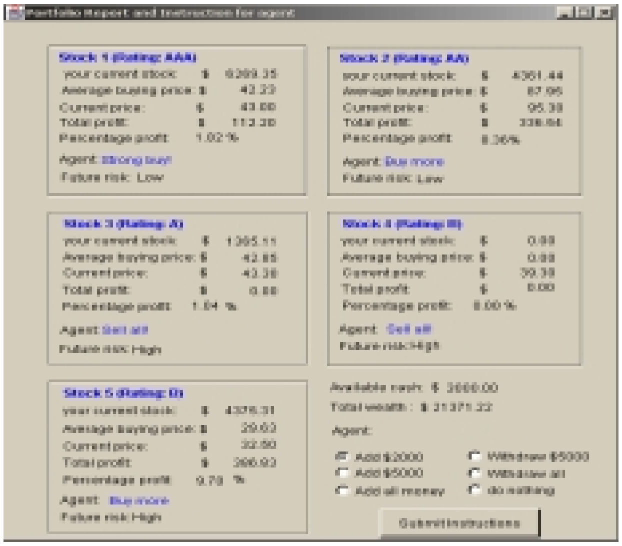

**A snapshot of the portfolio report in Experiment 2**

**A snapshot for market information in Experiment 2**

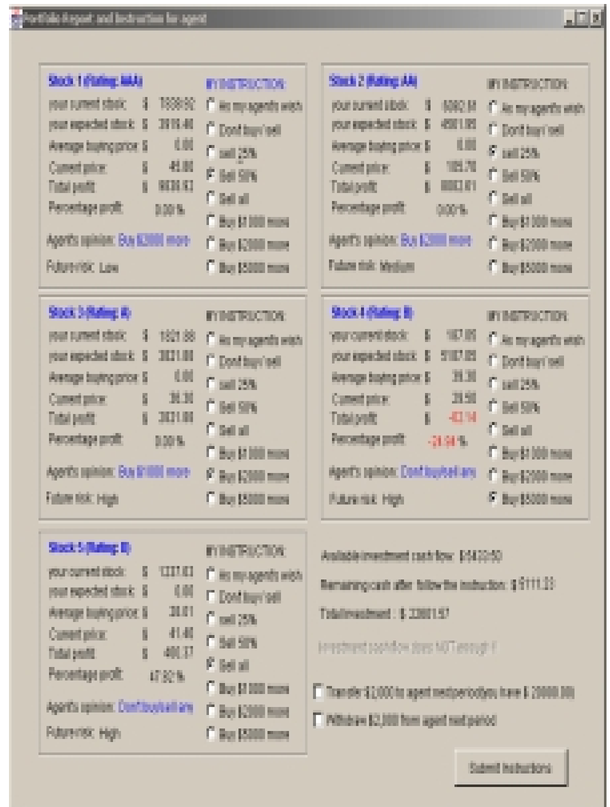

**A snapshot for the portfolio report in Experiment 3**

**A snapshot for market information in Experiment 3**