

The Role of Trust and Security Mechanisms in an Agent-Based Peer Help System.

Mike Winter
University of Saskatchewan
Saskatoon, SK Canada
mfw127@mail.usask.ca

Abstract

In this paper, we hope to document the use of various agent-based trust and security mechanisms in a simulated economic environment that is used to efficiently locate human helpers in a distributed peer-help environment. The agents and/or mechanisms will be responsible for ensuring that fraudulent or valueless transactions are monitored, and, if possible, the victim compensated. A variety of methods will be used further to this end, including: reference/review agents, police agents and relationship networks. The use of such agents or techniques will also be relevant to similar systems in the area of electronic-commerce or other agent-based applications.

The I-Help System

Introduction and Goals

The I-Help System (Greer et. al, 1998) is a distributed peer-help environment that uses an agent-based infrastructure to distribute help and advice in an efficient way. Each user of the system has a personal agent who acts on the user's behalf within the system. The agent keeps a model of its user, and uses the model to help communicate with other agents or a centralized 'matchmaker agent' when responding to requests from the user to either find or offer help. Apart from being an effective way to run a distributed peer-help environment, the use of agent technology as the foundation of the system allows us to experiment with and observe the emergent behavior exhibited by a society of agents each pursuing their own individualistic goals. These types of emergent behaviors will have global relevance across the domain of distributed agent-based systems, including electronic-commerce applications.

Underlying Economic System

As described in (Kostuik & Vassileva, 1999), a recent innovation added to the I-Help system is an underlying free-market economic system. By viewing advice and help as economic goods, and adding an artificial currency to the system, the distributed-help environment is transformed into a simulated economy with the goal of using emergent economic forces to distribute the help goods more effectively. Under this economic model, helpers and helpees would bargain through their agents using a simulated currency that could be traded in exchange for help. This approach allows an interesting solution to several dilemmas encountered in the basic I-Help system. For example, the system with the economic model provides extra incentive to the best tutors and students to participate, because they would no longer face excessive demand on their time due to efficient pricing and distribution mechanisms, and the simulated currency provides them with a tangible benefit for participation. A more detailed explanation of the system and its effects can be found in the above paper.

Similarities Between the I-Help System and E-Commerce Models

By extending the I-Help system to use an underlying economic model, the I-Help environment can be viewed as being analogous to other agent-based economic systems, such as those that exist in the e-commerce domain such as the OnSale or Ebay online marketplaces. While a distributed peer-help system for educational use does not provide a perfect mapping from the domain of distributed-help to the domain of e-commerce, relevant conclusions can still be made regarding agent behavior in electronic-commerce from experiments performed on the I-Help system. In the next section, we will examine how conclusions drawn from the I-Help system in regard to agent trust, reputation and security issues can be applied to the area of real-world e-commerce systems using a similar agent architecture.

Trust and Security in Online Commerce and Agent Societies.

(Zacharia et. al 1999) identify three major trust and security concerns in online marketplaces:

1. The first is that potential buyer of goods has no way of verifying the authenticity or quality of the good they are about to purchase. This is relevant both in an electronic-commerce situation where bidders, at best, can only rely on a grainy picture of the product they are to purchase, and in an educational environment, where a helpee only has the helper's word that they are capable of advising them.
2. The second is that either party in a transaction may not perform their contractual obligations, and due to the ephemeral nature of the online marketplace, enforcement mechanisms are either non-existent or ineffective. In an e-commerce environment, this problem is exhibited every time a buyer does not pay for goods received or a seller sends an inferior product to the one specified in the contract if they bother sending anything at all. In an educational environment, the helper may disconnect the help session immediately after agreeing to help, or some other form of contractual violation. Conversely, the helpee may demand help from the helper that goes beyond what is reasonably specified by their contract.
3. The third is that the identity of participants in the online marketplace may be unclear or unverifiable. The most effective trust and security mechanisms are useless when thieves, frauds or incompetents have the means to change their online personae. In an online marketplace, a party known to be a poor trader often finds it easier to reappear under a new name than attempt to try and improve their sullied reputation. This is an equal problem for both commercial and tutoring environments, though of course the consequences of such fraudulent behavior is likely to be more dire in a commercial setting.

The use of agents in online marketplaces allows the use of agent-based trust and security mechanisms to deal with these two issues. We will give an overview of some methods proposed to increase trust and security in online agent-based marketplaces, and analyze their applicability to the I-Help system and distributed-help systems in general.

Online Agent-Based Marketplace Trust, Control and Security Mechanisms

As described in (Rasmusson and Janson 1996), control mechanisms in online marketplace systems can be described in terms of 'soft' and 'hard' systems. Soft systems are identified by their more passive approach to security and trust issues, not relying as much on explicit punishment and exclusion as on market controls, trust mechanisms and references. Hard control, contrarily, relies much more on punitive and explicit punishments such as banning certain participants for the marketplace, refunds or fines.

Examples of soft control mechanisms include:

1. Centralized Reputation System

Using this sort of system, participants in an online marketplace rely on a centralized database to evaluate the trustworthiness of potential trading partners. The user or his personal agent can then take that information and include it in their utility calculation. If the economic actor sees an attractively priced good, but the reputation of the person offering it for sale is poor, they may choose to purchase a similar item at a higher price, but from a seller who has a better rating. Such a system parallels the real-world economic process of either consulting an organization such as the Better Business Bureau or a handyman's cooperative for advice on who to turn to.

Such systems are widely used in the world of online commerce. Ebay (Zacharia et. al, 1999) has a globally accessible reputation board for each member of their system, which is kept up-to-date for each particular user by the people he or she has dealt with. A positive review on a transaction earns the user a score of positive one, a negative review earns the user a score of negative one, and a neutral review leaves their score unchanged. The user's score, along with any comments left by people they have dealt with, are accessible to all members of Ebay.

Such a system is currently used by the I-Help system, where a centralized 'matchmaker' agent uses its own data, generally gathered by exit surveys taken after a help session is ended, along with the agents user models to give a list of recommended list of helpers to an agent who has made a help request. While this system has proved relatively effective in the I-Help system, it has several drawbacks. The first is that monopoly situations can easily arise when a centralized system is consistently making similar recommendations to all participants, which is disadvantageous in terms of efficient help distribution, scalability, individualization of help, and the overuse of certain helpers. Also, with this sort of system, there is no remedy for individual cases of fraudulent or incompetent behavior.

2. Decentralized Reputation Agents

Perhaps a more suitable mechanism for the distributed and decentralized nature of agent societies are decentralized reputation agents, whom themselves are actors in the economic sphere. Such agents are generally private services, and provide reputation information for a fee to interested clients. Examples in real-life include credit service companies and firms that provide investigations and background checks for a fee. Such systems do not appear to be in use in online marketplaces yet, but the potential for their use is high, due to the highly open and decentralized nature of Internet commerce. One area in both commercial and educational applications that would be well-suited to this sort of trust mechanism would be in specialized domains where a smaller, more specialized reputation agent would have an advantage over a larger, generalized and centralized reputation broker. Disadvantages of such systems include the additional overhead and complexity that would be added to a online marketplace featuring such agents. Other disadvantages include possible collusion between reputation agents and the agents they are ostensibly objectively

reviewing, and the problem of recursive reputation, i.e. do you have a reputation agent for reputation agents? Such complications probably render decentralized reputation agents ineffective for distributed peer-help systems.

3. Trust Networks and Collaborative Filtering

The two methods looked at in this section are all similar in that they are highly personalized and put more account in the subjective ratings provided by the user or his agent, thus leading to better individualization of trust ratings.

Trust networks are best illustrated by the Histos system described in (Zacharia et. al, 1999). In this system users in an online market rate each other based on their experiences in conducting transactions. The Histos system then maps all users and their ratings into a directed graph in which the nodes represent traders and the weighted arcs represent the trust rating between them. Then, the trust rating between two any users is a function dependent on the length of the path between them. It is then possible to do a breath-first search on the resultant graph and assign personalized trustworthiness ratings to each market option. Such a system appears to be quite a powerful tool for commercial purposes, and also holds potential for use in educational environments. Its disadvantages include the need for a pre-existing dense reputation tree before the system can become effective. However, the use of such a system in an peer-help environment provides an interesting solution to this problem; as the users of such a system will often have pre-existing informal real-life trust networks with their group of peers that could be installed into the network after collecting 'trust surveys' from the users of the system

Other approaches to modelling implicitly trust in recommendation systems and information filtering can be classified under the "soft" control category too. For example, user 'cliques' (Goldberg et al., 1992) exploit the fact that people who have shown a similarity in terms of their past needs and preferences can be grouped together and their future needs and preferences accurately predicted based on those of the other members of the clique. Collaborative rating as performed in Firefly (<http://www.firefly.net/>) system is an example of a collaborative-rating system used for e-commerce purposes.

'Soft' control agents and mechanisms are often most effective when used in larger and more open online marketplaces where the three primary security concerns stated above are more prevalent. In a distributed peer-help system such as I-Help, the authenticity of both the good being traded and the users doing the trading is much easier to monitor and regulate. Also, distributed peer-help systems, by definition, are smaller and the access to them more restricted. For this reason, the use of 'hard' control agents or mechanisms might be more feasible for distributed peer-help environments. Some examples of 'hard' social control that might be done manually, automatically, or by an agent, include:

- Refunds

If a party or parties in a transaction feel they have been cheated, a refund system would allow the transaction to be annulled and the money exchanged to be returned. This sort of system is not always feasible in online marketplaces because of the problem of actually finding out the facts regarding the transaction, establishing the identities of the parties involved, and making consistent decisions regarding refunds. In a distributed peer-help system where the help sessions are logged, the process would be easier, although the human overhead necessary for reviewing all disputed transactions might become prohibitive, making an agent-based solution more desirable.

- Expulsion

If a participant in an online marketplace is consistently acting in a fraudulent or incompetent manner, there could exist mechanisms which expel the participant from the online marketplace. The expulsion could be permanent or temporary, depending on the severity of the transgression. Expulsion could be handled by a 'police agent' acting under human supervision.

- Metering

One interesting method of hard control that may be only applicable to an environment similar to I-Help is the concept of 'metering', much like a taxi cab system. The two participants in a transaction would agree to a set rate before a help session takes place, the price paid being contingent on how long the transaction took. If the buyer was dissatisfied with the quality of help, he or she could cancel the session thus limiting the monetary loss. This system has the added advantage of compensating helpers if the session goes longer than they expected or wanted. This would be very advantageous in a system such as I-Help, where a major concern is the amount of tutoring load that falls upon the best students, and is also an interesting pricing option for similar applications

Simulation and Experiments

A simulation, using the SWARM simulation system was set up to model the I-help system with the distributed economic model. The SWARM system was chosen for its proficiency in modeling distributed systems. The economic engine of the system was a hybrid of the work done by (Steiglitz et. al 1996) and the work previously cited by (Kostiuk and Vassileva, 1999). There were several modifications made to the traditional economic model as to better mirror the unconventional nature of a peer-help economy:

- The "GDP" of the economy, which consists of the sum of help value, or knowledge, contained by the individual actors in the system, never goes down. Unlike in a real economic system, where the supply of goods can fluctuate without bound, we wished to capture the reality that the total of knowledge in most peer-help environments, especially in educational settings, will never decrease over time, but rise as knowledge is shared within the system.
- Although a marketplace is usually very decentralized, we instituted a centralized "matchmaker" system that handles all bids and reputation checking. This is done both to properly model the I-Help system as it actually exists, and for the sake of convenience in instituting control mechanisms. This is a realistic step, as most distributed peer-help environments will be collated around a central organization that will provide the peer-help infrastructure for the entire group. The matchmaker matches agents up by assigning each helpee to the helper whose price for help is the closest bid of equal or lesser value to the bid of the helpee. This is changed when a reputation system is implemented, wherein the matchmaker matches the helpee with the helper whose bid is of equal or lesser value and whose reputation value is the highest.

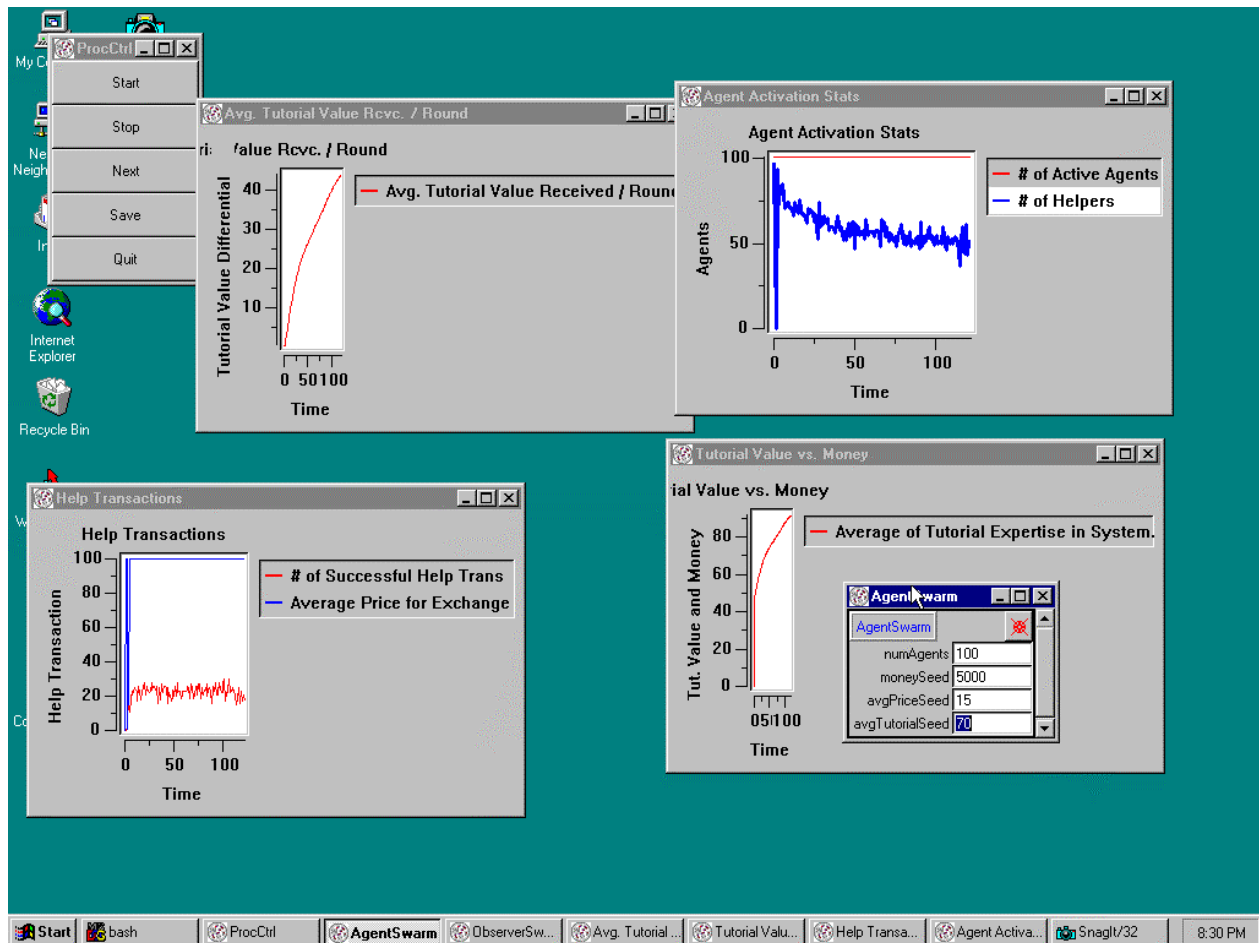


Fig. 1 A screenshot of the simulation.

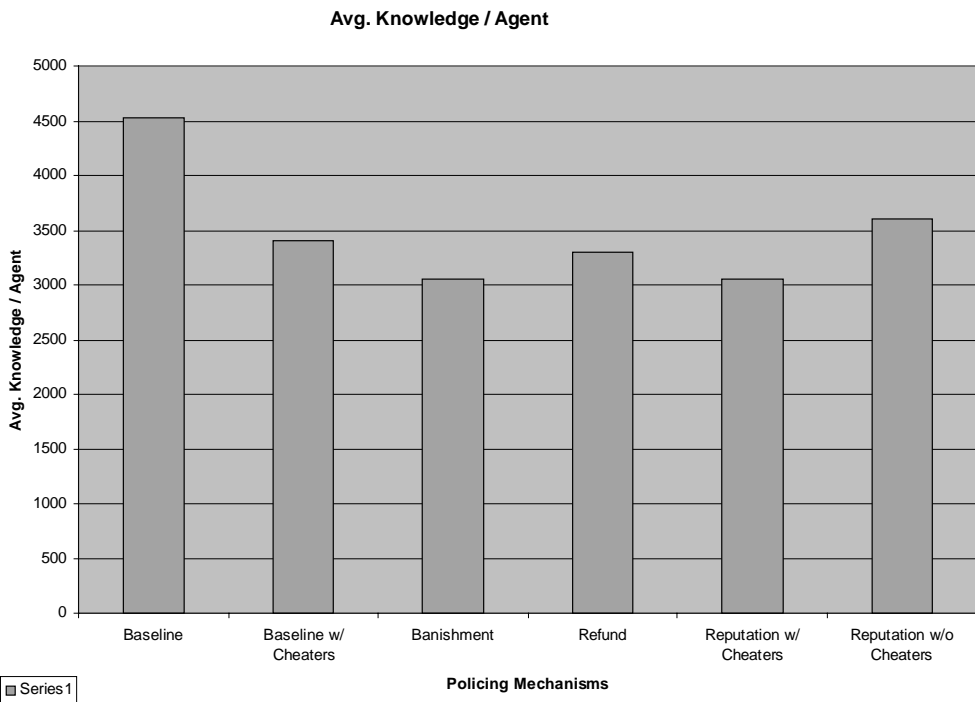

- Another interesting modification to the simulated economy is that when a helper helps someone, they increase their knowledge as well. This follows from the fact that tutors usually gain knowledge of a subject when they teach it. We therefore have helpers increasing their knowledge by a fixed fraction of how much help they give in each help transaction.
- The utility function for each agent is the algorithm that determines whether or not the agent enters the market as a buyer or seller and how much of the help good it wants to buy or sell. The utility algorithm used in the simulation takes into account the relation of the agent's knowledge relative to the mean knowledge of all agents in the simulation. This is meant to model the educational situation in which poor students will more likely be concerned about increasing their marks than increasing their money, and the better students will be more interested in money than in increasing their mark.
- A price ceiling. To achieve all of its goals, a peer-help system must avoid hyperinflationary market situations. This is important for several reasons: a hyperinflationary system will drive people away from using the system because they do not know the worth of their virtual money or

their knowledge, wild fluctuations often cause the underlying economic system to flatline or diverge to infinity, and volatile systems are prime targets for fraud, speculation and cheating.

With these modifications to a normal economic simulation implemented, we constructed a preliminary simulation to compare and contrast various policing and trust mechanisms in a distributed peer-help environment. At the time of this paper, three of the mechanisms outlined above have been tested, and there are plans to test at least two more before implementation of the new I-Help system. It was decided that the following measures would be the most useful in determining which mechanism was most effective:

- i) The average amount of knowledge per agent at the end of the simulation
 - This is the most important metric. The purpose of a peer-help system is to increase the amount of knowledge in as many people as it can.
- ii) The number of transactions per round.
 - Another goal of peer-help systems is to have as many people involved in them as possible. An active system is a healthy system.
- iii) The amount of knowledge gained per round by each agent.
 - This metric is a good measure of the quality of the transactions in each round. A high number means there were most likely many good transactions in that round.

Six tests were performed. The first involved testing an optimal system with no deception involved. The second involved testing a system wherein every twelfth agent was a “cheater”, a class of agents that would routinely overcharge for the amount of tutorial knowledge they had to offer. The third involved testing a system in which the cheating agents were discovered just before the halfway mark of the simulation run-time, and barred from participating any further. The fourth involved testing a mechanism in which refunds were given to agents who received help that had value far below what they paid for it. The fifth and sixth tests involved the implementation of a centralized reputation system as described above, and it was tested with and without the presence of cheating agents. The results of these tests is graphed and summarized below. In the “average tutorial per agent metric” (fig. 2) we see that the straight free market system works the best, with the baseline reputation system coming in second. Surprisingly, the baseline economic system with cheaters added ended up outscoring all of the security mechanisms that were put in to neutralize the effect of the cheaters, suggesting that a libertarian, free-market system might be effective, or that the simulation has some odd dynamics which not might map to the real world. The refund mechanism scored the best of the policing measures. The agent involvement numbers (fig. 3) are fairly intuitive. The highest scorers are the baseline systems, and the refund system, which appears to do a very good job at keeping the market active. The worst scorer was the banishment system, which is to be expected



Figs. 2 and 3

The “quality of help”(fig. 4) indicator is where we see the most pronounced influence of the policing mechanisms. The two reputation mechanisms and the banishment mechanism score the highest on this metric.

Overall, in this preliminary example, there seems to be no clear winner as to which policing mechanism is the most effective for all goals, although a mechanism which could take advantage of both an efficient pricing system and reputation data would be very powerful. Further testing of options should help resolve the situation.




Fig. 4

Conclusion and Further Work

The tentative results gleaned from the simulation suggest a positive result to the experiment of using a peer-help system with an economic distribution system using trust and policing mechanisms. Along with data from the simulation, the use of the I-Help system with the underlying economic system is tentatively planned for use in a class of a hundred students in the 1999 fall semester. Using this data, we will be able to determine which categories of online marketplace security controls are best for a distributed peer-help environment, and the same data will be valuable in determining what sort of agent-based control mechanisms will be effective in a e-commerce environment.

Further work involves refinement of the simulation, including bringing it further into line with the theoretical work described in (Kostiuk and Vassileva,1999), the testing of additional trust and

policing mechanisms both separately and in combination, and the use of empirical data garnered from the classroom tests to both refine the I-Help system for use in its current domain and to further the use of this novel system into other situations where a distributed, agent-based peer-help system would be of use.

Acknowledgements

I would like to acknowledge the contribution of my supervisor, Julita Vassileva, for guiding my research into this subject. I would also like to thank Kevin Kostiuk for helping with the design of the simulation, and for sources of research and inspiration and Vive Kumar for help in formatting the paper. Finally, I would like to thank Jim Greer, Ralph Deters and the rest of the Telelearning group for funding, ideas and discussion. Funding for this paper was provided in part by the Telelearning-NCE Project.

References

Greer J., McCalla G., Cooke, J., Collins, J., Kumar, V., Bishop, A. and Vassileva J. (1998) The Intelligent Help Desk; Supporting Peer Help in a University Course, Proceedings ITS '98, San Antonio, Texas.

Goldberg D., Nichols, D., Oki, B., and Terry, D. (1992) Using collaborative filtering to weave the information tapestry. Communications of the ACM 36 (12), 61-70.

Kostiuk, K. and Vassileva, J. (1999) Free Market Control for a Multi-Agent Based Peer Help Environment, Proceedings of the Workshop on Agent-Based Decision Support for Managing the Internet-Enabled Supply Chain, Agents 99, Seattle WA.

Maes P. "Agents that Reduce Work and Information Overload." Communications of the ACM 37 (7), 31-40, ACM Press, July 1994.

Rasmusson, L. and Janson, S. (1996) Simulated Social Control for Secure Internet Commerce, New Security Paradigms '96, ACM Press, Sept. 1996.

Steiglitz, Ken, Honig, Michael L. and Cohen, Leonard M. "A Computational Market Model Based on Individual Action", Market-Based Control: A Paradigm for Distributed Resource Allocation; ed. Scott H. Clearwater. World Scientific; 1996, River Edge, N.J.

Zacharia, J., A. Moukas and Maes, P. (1999) Collaborative Reputation Mechanisms in Electronic Marketplaces, Proceedings of Electronic Commerce Technology Minitrack, HICCS-32, Wailea Maui, Hawaii