

Detecting and Preventing Defamation in Multiagent Systems

Elijah Bitting, Jonathan Carter, and Ali A. Ghorbani
Faculty of Computer Science
University of New Brunswick
Fredericton, NB, Canada

ABSTRACT

A new theoretical model for the control of defamation in Multiagent Systems (MAS) is proposed. The goal of the proposed anti-defamation model is to develop effective mechanisms to protect all types of systems from defamation. The proposed model differs from a more general deception handling model in that only data relating directly to the calculation of reputations are evaluated for deception. An e-commerce simulation based on this model is described and results are discussed. Results from the experimental simulations support the validity of the proposed anti-defamation model, as they show the anti-defamation model is successful in protecting the simulated MAS from the overall damage caused by defamation. A concerning outcome however, is that while all agents are better off with the proposed anti-defamation model, the model needs further improvements in order to be successful in targeting and punishing specific dishonest agents. Although this work is directed toward defamation in particular, we believe it may be of particular interest to anyone wishing to provide protection against all types of deception.

Categories and Subject Descriptors

I.2.11 [Computing Methodologies]: Artificial Intelligence—*Multiagent Systems*; I.6.5 [Computing Methodologies]: Simulation and Modeling—*Model Development*

General Terms

Design, Performance, Experimentation, Legal Aspects

Keywords

Defamation, deception, suspicion, trust, multiagent systems

1. INTRODUCTION

Defamatory speech, along with many other types of mendacious speech [10], are pervasive in various areas of our society, from personal to corporate relationships, causing great

harm to innocent individuals and groups. The issue of reputation damage is largely ignored in the study of Multiagent Systems (MAS).

Effective mechanisms which allow agents in a MAS to reason about the others' reputations invariably require that all (or at least some) of the agents report specifics about their interactions with other agents to the rest of the society. This reporting allows the society as a whole to gain insight into how specific agents might act in future interactions. In other words, the feedback itself is one of several components necessary to build reputation. Other sources also describe trust mechanisms that rely on the ability of agents to provide some similar sort of objective social feedback [2, 16].

When a subjective feedback mechanism is integrated into a MAS, the agents are given the freedom to do what they wish with this feedback. This freedom provides the agents with the power to indirectly manipulate the opinions of others. Left unchecked, this power could have dire consequences.

To the best of our knowledge at the time of this writing, there does not exist any publicly available work which explicitly addresses the issues mentioned above, concerning defamation of any sort in MAS. As a result of this work being somewhat new in the world of MAS, relatively few published resources have been found to compare to the ideas presented here. Some work has been done regarding the control of defamatory behavior in online reputation systems, but the resulting methods are based on controlled anonymity, which may not always be viable depending on the situation. Using any type of anonymity mechanism is not effective when the type of interaction inherently identifies one or both parties involved. Such interactions include stays at hotels or encounters with restaurants or doctors [2].

2. RELATED SOCIOLOGICAL ISSUES

This section contains discussion relating to several concepts closely tied to defamation. Each is important for a full understanding of what defamation is and how it may be controlled in MAS.

2.1 Trust

The hope of integrating trust mechanisms into MAS is that by making use of such facilities, social interaction will be facilitated in the face of uncertainty. In addition, and more importantly for this work, defamation protection mechanisms are intended to help maintain stability in such 'trusting' systems. This stability results from protection against unfair attacks on individuals' reputations.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CNSR 2003 Conference, May 15-16, 2003, Moncton, New Brunswick, Canada.

Copyright 2003 CNSR Project 1-55131-080-5 ...\$5.00

As noted by Hwang in [6] trust is a key concept in most interpersonal and group behavior. An inherently vague and multifaceted concept, trust becomes especially cloudy when attempting to define it formally. There are numerous varying definitions of trust present in the literature. The reason for the apparent disagreement is inherent in trust. Trust is domain specific, so its definition is dependent on the specific context in which one is working. Trust is also subjective as it not only depends on context but more specifically on the beholder.

2.2 Reputation

While held by the individual, reputation is ascribed by society as a whole and cannot be established by the individual to whom the reputation belongs. The society determines reputation through a general estimation of the individual in question; dependent on the individual's management of their own identity and how they present this identity to the rest of society. Through this presented self, society can evaluate reputation based on the objective behavior of the individual in question [14].

The core of reputation is the role or expectation. Reputation is the assessment by society (or some segment thereof) of an individual's fulfilment of that society's expected roles. These roles are a central part of any society and the production thereof depends on the collective beliefs and specific context of the society. In addition, roles established for the evaluation of reputation are necessarily linked to the society's goals. In particular, any such roles established in the society must reward actions which help to achieve the society's goals.

Reputation is ascribed by the society to which an individual belongs, and not by the individual itself. Thus, an individual is incapable of directly manipulating their own reputation. Through specific manipulation, however, individuals can manipulate and in some cases deceive others into ascribing higher reputations than they would otherwise.

As suggested above, it is from this realm of reputation that defamation emerges. Only when we establish reputation in a society is it possible for damage to be caused to that reputation. One cannot defame another if there is no concept of reputation in the society. In particular, reputation building mechanisms that allow members of society to subjectively report on others, open the door for untrue reports. The effect of these untrue reports is usually consistent with the desired effect: damage to an individual's reputation. It is this case which is of interest for this paper. When an individual's reputation is damaged as a result of untrue statements made by another to a third party, these untrue statements are classified as defamatory.

2.3 Defamation

To start off the discussion on defamation itself, a formal definition of defamation will be given followed by a few related legal notes. As defamation is a social concept, it is defined in terms of the nature of social communication.

2.3.1 Defining Defamation

We define defamation as a deceitful communication which unfairly harms the reputation of the individual or group being lied about. The statements must be *harmful* to the individual claiming to be the victim of defamation. A defamatory statement must be made to a third party, meaning

that statements made directly to the individual who is the subject of those statements cannot be considered defamatory because they could not possibly effectively harm the individual's reputation. To illustrate, consider two acquaintances locked in a sealed room. As long as another does not overhear their conversation and they speak only of each other, no defamation is possible'.

In order to prove a case of defamation in the courts, the plaintiff must prove that the statements in question were false, caused actual damage to their social reputation and (as already noted above), that the statements reached a third party.

The most obvious challenge when it comes to laws regarding defamation is that they are often inapplicable due to defenses based on various forms of freedom of expression [9].

2.4 Deceit

Deceit involves knowingly giving information to another which is inaccurate, or otherwise untrue.

2.4.1 Utility of Deceit

In [5], a game is described where agents can greatly benefit from deceiving their opponents. Several other papers also exist that discuss deceit from various related perspectives: trust and security in MAS [15], deception in agent negotiations [17] and internet security [13].

Deception ultimately renders the deceived party's decision making process useless, as the information they are working with is erroneous. Researchers have outlined many reasons for deceit, such as the desire to induce certain states of mind in their neighbors [1]. As stated above, an individual cannot directly manipulate their own social reputation. While this is true, using deception they can manipulate others into ascribing them a higher reputation than they would otherwise.

2.5 Suspicion

Suspicion is an important element in modelling defamation in MAS, because it is the key human mechanism used to alert individuals to the presence of defamation or other sorts of deception or wrongdoing.

Suspicion requires that agents reason in an environment of uncertainty [1, 7]; loosely analogous to environments with partial information. A model for reasoning in such an environment of uncertainty is presented in [1], where *plausibility* and *credibility* are measured in order to establish a suspicion value. Consider for example that an information source, S (sender), sends a fact, p , to a receiver R . R 's suspicion in S 's p depends on: 1) Plausibility that S actually believes p and 2) S 's convenience in inducing R to believe p . Plausibility is covered in detail in [1], including methods for calculation.

2.5.1 Opinion Model

We make use of the work of Josang and Knapkog [7], who have developed a pair of models for reasoning in uncertain environments.

The Evidence Space model [7] is used to predict the probability of a binary phenomenon by performing calculations based on the number of observed positive and negative results (r and s respectively).

Related to the evidence space by Equations 1, 2, and 3, the opinion space model involves building opinions about facts or events in the form $\omega_p = \{b, d, u\}$ where $b + d + u = 1$. ω_p is referred to as an opinion about some binary event p and b ,

d and u are the measures of belief, disbelief and uncertainty, respectively.

$$b = \frac{r_p}{r_p + s_p + 1} \quad (1)$$

$$d = \frac{s_p}{r_p + s_p + 1} \quad (2)$$

$$u = \frac{1}{r_p + s_p + 1} \quad (3)$$

3. ANTI-DEFAMATION MODEL

What is proposed is that agents be equipped with a reasoning mechanism to allow them to become suspicious of inconsistent information relating to reputations. This suspicion can then be used as a trigger to request a consensus in the community. The consensus, while not always exact, can be used as an estimate to the sincerity of the party with whom the agent is interacting. The results of this suspicion triggered consensus is then recorded as a new opinion by incrementing the appropriate counter from the opinion model.

We call this mechanism a defamation module, shown in Figure 1. The defamation module receives reputational evaluations of each agent with which it interacts. This reputational data results from individual observations as well as feedback and recommendations from others.

3.1 Reputation

In this model the reputation of an agent is determined by its fulfilment of Only a single role. This role is called the competitive role. The competitive role is defined in terms of relative pricing; relative to the other agent's prices. The magnitude of satisfaction of this role by agent q , as seen by agent p is denoted C_q^p .

The reputation of an agent q as seen by agent p is denoted R_q^p .

3.2 Suspicion

Agents are said to become suspicious of defamation when their view of one or more agents' (other than their own) reputations are reduced to a significant enough degree as a result of some communication. The reputation change can be easily calculated if the agent constructs a hypothetical view of the other agents in the MAS. Hypothetical means that the agent assumes the received communication is true and constructs a new artificial set of reputations. The reputation change values are easily obtained by subtracting the new hypothetical reputation values from the real established reputations.

Standard deviation is used as it is a well known concept for analyzing risk in various areas such as economics [12] and project scheduling [4]. In this case the standard deviation of the perceived reputations of sellers is a measure of the variance amongst the buying agent's view of the seller's reputations.

3.3 Consensus

The consensus algorithm used here involves an agent sending a consensus request message to all other agents in the MAS, and waiting for enough responses to ensure the majority are truthful. This is made possible by adopting the constraints defined below.

To make the consensus algorithm used in this model more efficient, two constraints are placed upon the model itself:

- The maximum number of potential liars must be controlled and known. More precisely, for the consensus algorithm to work effectively, the maximum possible number of liars, m , is restricted by $2m < n$ where n is the total number of agents.
- When waiting for responses to a consensus request, agents should only wait for the first $2m + 1$ responses. Any additional responses are superfluous.

This method of restricting the maximum number of liars to m and waiting for only $2m + 1$ responses allows for many fewer messages to be sent in reaching a consensus, particularly in large systems where $m \ll n$.

3.4 Opinions and Certainty

The results of the above-described suspicion and consensus evaluation are recorded as estimations of truth and untruth. Each agent maintains two sets of counters, r_i and s_i . Every agent holds both r_i and s_i associated with every other agent i . r and s are the positive (truth) and negative (untruth) counters respectively. The resulting opinion (ω_i) and its disbelief and uncertainty elements (d and u) are used to qualify the agents reputation. When working with reputations the actual reputation value \hat{R} is used instead of R , where \hat{R} is calculated as follows:

$$\hat{R} = R - \max(d - u, 0) \quad (4)$$

4. DEFAMATION SIMULATION

The simulation consists of two types of agent: buyers and sellers. Buying agents are dispatched to selling agents after calculating the shopping factor, δ (defined below). The seller quotes their price along with the prices of their competitors. Upon receipt of the price quote, the buying agent will compute a set of hypothetical reputations that would result if the price quote were truthful. These hypothetical reputations are then compared to the currently held reputations. Through this comparison, the buying agent determines if the quote is suspicious. If the agent suspects a price quote, a consensus is requested. If a majority of the consensus responses indicate a truthful quote, the buying agent continues as if no suspicion had occurred. If on the other hand the consensus results suggest deceit, the buyer records the occurrence of a lie authored by the selling agent in their 'opinion'. A lie is recorded as an increment to the s (negative) parameter to the evidence space.

The selling agents fix their prices in a supply/demand fashion. While active, if the seller agent sells more than one unit on an iteration, they increase their price by \$1. If they had one or no sales on the last iteration, they maintain the same price. After 10 consecutive iterations with no sales, a seller agent is said to become inactive. When inactive, an agent lowers its price by \$1 every five iterations. Once a seller has become inactive it must either sell to more than one buyer on a single iteration, or sell to at least one buyer on each of two consecutive iterations the agent becomes active again. While inactive, dishonest sellers cease to lie, they begin lying again immediately upon reactivation.

The shopping factor is calculated in order to establish a factor by which decisions can be made based on both trav-

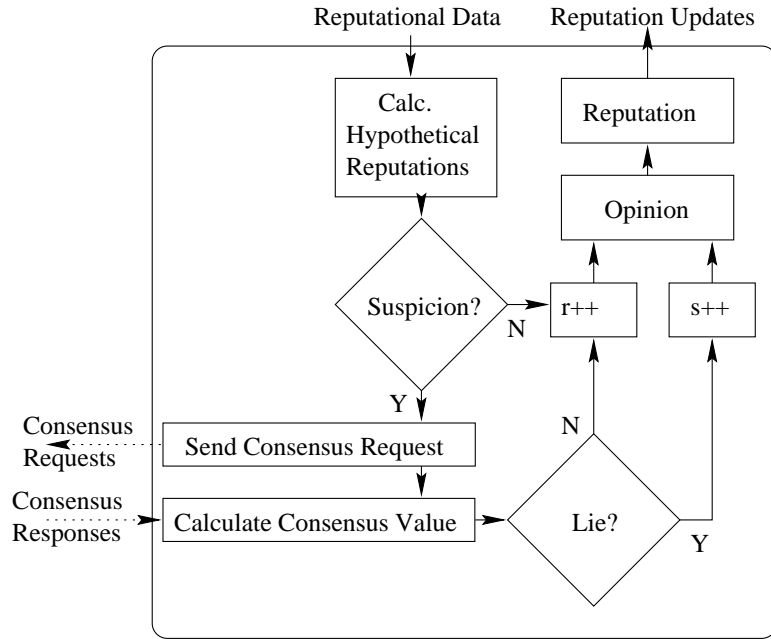


Figure 1: One Agent's Defamation Module

elling distance and reputation. The shopping factor calculated by buyer p , for seller q , is denoted δ_q^p and calculated as shown in Equation 6.

The game is played in a virtual cartesian-like two-dimensional neighborhood. Let D represent the maximum possible travelling distance in this neighborhood. Also let dd represent $f((x_1, y_1), (x_2, y_2))$ or the travelling distance. dd is calculated by summing the differences of the x and y coordinates of the two locations being considered. This is shown in Equation 5.

$$dd = f((x_1, y_1), (x_2, y_2)) = |x_1 - x_2| + |y_1 - y_2| \quad (5)$$

$$\delta_q^p = \frac{R_q^p}{\alpha(1 + \frac{dd}{D})} \quad (6)$$

Where the buyer is located at (x_1, y_1) and the seller at (x_2, y_2) . α is used as a parameter to control the extent to which the shopping factor is dominated by travelling distance, relative to reputation. For our simulation α was set at $1/2$.

On each game iteration, each buying agent calculates a shopping factor in regard to each selling agent. The most favorable (largest) factor tells the buying agent where to shop in the current iteration. This factor causes agents to shop at various locations depending on the distance to those locations and the reputations of the selling agents that reside there.

4.1 Parameters

- n_B - Number of buying agents.
- n_S - Number of selling Agents.
- m - Number of liars.

By lying about its competitors prices, the selling agent is defaming its competitors. First of all, the price quotes being discussed are false, this is requirement number one for defamation. Secondly, if believed, the false quotes are damaging to others reputations. Third and finally, the untrue price quotes are sent to buyers who must be a third party. Thus, these false price quotes constitute defamation in the context of this MAS simulation.

5. ANALYSIS

Three performance measures are used in analyzing the results from these simulations. Cumulative sales, sales per iteration and activity rates. For all three of these measures, the group of honest and dishonest sellers are considered separately for comparison purposes.

Three distinct cases are examined:

- Constant sellers and liars, varying buyers.
- Constant buyers, varying sellers, varying liars but liars remains a constant proportion of sellers.
- Constant buyers and sellers, varying liars.

These simulations are named *SimA*, *SimB* and *SimC* respectively.

In addition to the above described simulations, a set of control experiments are conducted with identical parameters, only without anti-defamation mechanisms enabled. These simulations are named *SimA*, *SimB* and *SimC*.

The *SimA* simulations are characterized as follows: 50 sellers, 10 of which are liars, buyers vary in *SimA₁* to *SimA₇* as follows: 5, 25, 50, 75, 100, 125, and 150.

The *SimB* simulations are characterized as follows: constant 50 buyers, constant proportion of sellers are liars ($1/2$), and number of sellers for *SimB₁* through *SimB₈* are 3, 25, 50, 75, 100, 125, 150, and 175 respectively.




Figure 2: Cumulative Sales for both \overline{SimA} and \widehat{SimA}

The $SimC$ simulations have a constant 50 buyers and 50 sellers, but the number of dishonest sellers changes from $SimC_1$ to $SimC_7$ as follows: 0, 1, 10, 20, 30, 40, and 50.

When the averaged results from the $SimA$ simulations, \overline{SimA} , are compared to the corresponding control simulations without anti-defamation (\widehat{SimA}) improved sale conditions for all selling agents may be observed, including both the honest and dishonest agents. Figure 2 illustrates this result. This figure illustrates how the honest and dishonest agents both fair better with the anti-defamation model.

In order to examine the performance of honest versus dishonest sellers, the growth rates of dishonest and honest agents' sales curves are compared. Growth rates corresponding to $SimB_2$, $SimB_3$, $SimB_4$ and $SimB_5$ are shown in Figure 3, and labelled a , b , c and d , respectively.

As the number of sellers increases the growth rates of both honest and dishonest agents smooth until they are nearly static in $SimB_4$. Although in $SimB_2$, the honest sellers do begin with a larger sales growth rate than the dishonest sellers. That rate is eventually overtaken by that of the dishonest sellers. This trend does not continue however; after one simulation with honest sellers maintaining a slightly higher sales growth rate (c in Figure 3) the dishonest sales growth rates settle above the corresponding honest rates. An equilibrium is approached as the number of selling agents exceeds the number of buyers. While it would be desirable for the dishonest agents to be forced to receive lower sales than the honest ones, this result is still encouraging as it does show stability resulting from use of the proposed anti-defamation model. In addition, although the dishonest agents slightly outperform their honest counterparts, the outperformance is more severe, and less stability is observed when the anti-defamation model is not used.

Figure 4 shows results similar to those shown in Figure 3, except for \widehat{SimB}_2 through \widehat{SimB}_5 instead of $SimB_2$ through $SimB_5$. Note that the two figures have different scales. Figure 4 shows the slopes of the cumulative sales curves for both honest and dishonest agents in the following simulation runs: \widehat{SimB}_2 , \widehat{SimB}_3 , \widehat{SimB}_4 and \widehat{SimB}_5 labelled as a , b , c and d respectively. Notice as you compare Figures 3

and 4, the pairs of curves representing honest and dishonest sale growth rates in Figure 4 do not converge to a steady pattern as those shown in Figure 3.




Figure 3: Slope of cumulative sales: $SimB_2(a)$, $SimB_3(b)$, $SimB_4(c)$, $SimB_5(d)$




Figure 4: Slope of cumulative sales: $\widehat{SimB}_2(a)$, $\widehat{SimB}_3(b)$, $\widehat{SimB}_4(c)$, $\widehat{SimB}_5(d)$

In $SimC$, the proportion of dishonest sellers seems to have little effect on the performance of the marketplace until that proportion increases over one half. When the ratio of dishonest agents to honest agents is low, the dishonest agents perform only slightly better than the honest agents in terms of sales. When the proportion of dishonest sellers exceeds 50%, the dishonest group begins to gain ground widening the gap between themselves and the honest sellers. Three graphics showing the honest and dishonest cumulative sales figures corresponding to $SimC_3$, $SimC_4$ and $SimC_5$ are shown in Figure 5 to illustrate these results. Note that the percentages of dishonest sellers for $SimC_3$, $SimC_4$ and $SimC_5$ are 20%, 40% and 60% respectively. In the third graphic corresponding to $SimC_5$ in Figure 5, the dishonest sellers begin

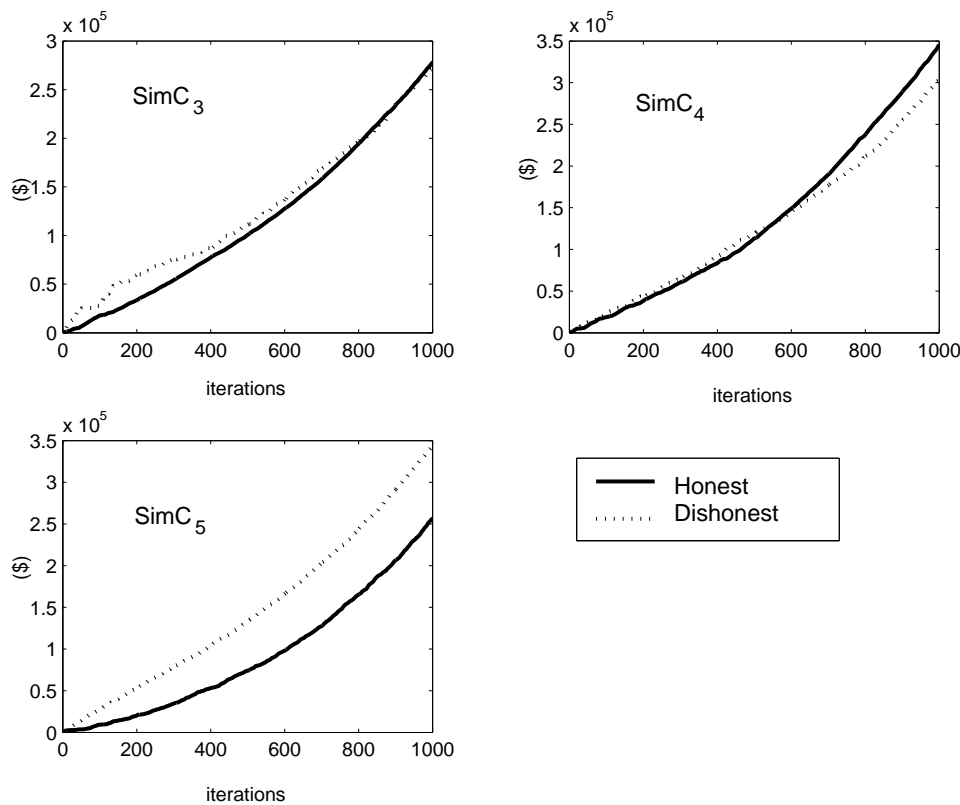


Figure 5: Cumulative sales from *SimC*₃, *SimC*₄, and *SimC*₅

to widen the gap between themselves and the honest sellers. A similar story is told for *SimC*, except that the dishonest agents always outperform the honest agents to a significant degree regardless of their proportion.

5.1 Scalability

The maximum number of selling agents we are able to include in these simulations is 150. This restriction is due to the memory demands of the simulation program in execution. With any more than 150 buying agents, the simulation executes for sometime, often reaching 500 – 600 iterations, then crashes with a memory error. The memory constraints are imposed by the systems upon which the simulations are run and not the complexity of the system itself. The game is scalable, given extended memory and processor capacity.

The number of selling agents is increased until the simulations will no longer run due to memory errors. In the case of 1000 iterations and 50 buying agents, the maximum number of selling agents simulated is 175.

Time complexity is not an issue for simulations of this size. Some of the larger simulation runs did take several minutes; but the program would crash due to storage complexity not lack of time.

5.2 Stability

For the game currently being discussed, we define stability as the convergence of sales growth rates to a consistent pattern. In comparison to Figure 4, Figure 3 shows a definite convergence of both honest and dishonest sales rates to near constant and near equivalent rates. Figure 4

shows no definite pattern of convergence. Recall that Figure 4 shows sales growth rates corresponding to simulations without anti-defamation mechanisms and Figure 3 shows the corresponding simulations with anti-defamation. This shows the stabilizing effect of the proposed anti-defamation model on the evaluation game being discussed.

6. CONCLUSIONS AND FUTURE WORK

Due to the ‘soft’ nature of the defamation response mechanism of the proposed model, instead of dishonest agents becoming ostracized in the community, the liars seem to be victims of only slight punishment. This fact is due largely to the design decision to allow the agents to have the capability to become inactive. While inactive, the dishonest agents stop lying and may re-establish reputation through reducing prices and providing trustworthy information. This need for dishonest agents to become inactive and re-establish damaged reputation, may provide a partial explanation to the fact that honest agents tend to maintain higher activity rates when compared to their dishonest counterparts.

Future work is intended to refine the proposed model. First is the need for a more multifaceted reputation mechanism. A better reputation mechanism would allow for more meaningful analysis of variations in reputation. In addition, it is desirable to construct more meaningful methods for determining suspicion as a results of variance in observed data. In particular, a non-binary suspicion mechanism which makes use of the opinion model was initially pursued but abandoned due to the simplicity of the reputation mechanism being used. Initially, an opinion was used to

model an agent's suspicion, but no meaningful cutoff point was available to determine the presence of a lie.

Due to the non-distributed synchronous nature of the current simulation and the limitations of a single processor windows environment, the program consumes great amounts of memory while executing. This has a negative impact on scalability. Future work is intended to extend this system to an asynchronous distributed MAS which will be far more flexible in terms of scalability.

Future work with this model will undoubtedly turn to the statistical work of Dempster and Shafer. Dempster and Shafer's work relates directly to the use of statistical evidence as a decision making tool. See the following sources for more information related to this work [3, 11, 8].

7. REFERENCES

- [1] F. de Rosis, C. Castelfranchi, and V. Carofiglio. On various sources of uncertainty in modeling suspicion and how to treat them. In *Deception, Fraud and Trust in Agent Societies*, pages 61–72, June 2000.
- [2] C. Dellarocas. Immunizing online reputation reporting systems against unfair ratings and discriminatory behavior. In *ACM Conference on Electronic Commerce*, pages 150–157, Minneapolis, MN, 2000.
- [3] A. P. Dempster. Upper and lower probabilities induced by multivalued mapping. *Annals of Mathematical Statistics*, 38:325–337, 1967.
- [4] G. L. D. Furia. Measure schedule risk using the standard deviation. *ESI Horizons*, 3(9), January 2002.
- [5] J. P. Hespanha, Y. S. Ateskan, and H. H. Kizilocak. Deception in non-cooperative games with partial information, tech. report, EE Systems, University of Southern California, Los Angeles, CA, April 2000.
- [6] P. Hwang and W. P. Burgers. Properties of trust: An analytical view. *Organizational Behavior and Human Decision Processes*, 69(1):67–73, January 1997.
- [7] A. Josang and S. Knapskog. A metric for trusted systems. In *National Institute of Standards and Technology National Computer Security Center 21st National Information Systems Security Conference*, Hyatt Regency Crystal City Arlington, Virginia, USA, 1998. Chapman & Hall.
- [8] M. Lalmas. Dempster-shafer's theory of evidence applied to structured documents: Modelling uncertainty. In *Proceedings of the 20th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 110–118, New York, 1997. ACM.
- [9] P. D. Lawler. Group defamation, submissions to the attorney general of ontario, March 1984.
- [10] F. Peonidis. Freedom of expression, autonomy, and defamation. *Law and Philosophy*, 17:1–17, 1998.
- [11] G. Shafer. *A mathematical theory of evidence*. Princeton University Press, Princeton and London, 1976.
- [12] D. R. Stabile and B. H. Putnam. Irving fisher and statistical approaches to risk. *Review of Financial Economics*, 11(3):191–203, 2002.
- [13] D. Toft. Ftc sets up to sniff out internet scams. *PC World Online*, URL: <http://www.pcworld.com/resource/printable/article/0,aid,11585,00.asp>, June 25 1999.
- [14] N. UK, H. Press, and G. Erving. *The Presentation of Self in Everyday Life*. Horizon Scientific Press, Norfolk, UK, 1959.
- [15] H. C. Wong and K. P. Sycara. Adding security and trust to multiagent systems. *Applied Artificial Intelligence*, 14(9):927–941, 2000.
- [16] G. Zacharia, A. Moukas, and P. Maes. Collaborative reputation mechanisms in electronic marketplaces. In *Proceedings of the 32nd Hawaii International Conference on System Sciences*, Hawaii, 1999.
- [17] G. Zlotkin and J. S. Rosenschein. Incomplete information and deception in multi-agent negotiations. In *In Proceedings of the twelfth international joint conference on Artificial Intelligence*, pages 225–231, Sydney, Australia, 1991.