

Towards a Formalization of Trust

Jonathan Carter and Ali A. Ghorbani

Faculty of Computer Science
University of New Brunswick
Fredericton, NB, E3B 5A3

Abstract. This work focuses on the design and implementation of a new model of trust. The new model of trust is based on the formalizations of reputation, self-esteem, and similarity within an agent. Our previous work establishes the formalization of reputation within an information-sharing Multiagent System. The previous work claims that reputation cannot be universalized. This work universalizes reputation through the use of values within all Multiagent Systems. The following values are shown to be manifested within Multiagent Systems: responsibility, honesty, independence, obedience, ambition, helpfulness, capability, knowledgeability, and cost-efficiency. Manifestations of these values result in a more universalized approach to formalizing reputation.

Self-esteem is formalized as the reputation an agent has with itself. Lastly, similarity is formalized as the difference in the importance of the values previously mentioned. Combined, the weighted components of self-esteem, similarity, and reputation form a new model of trust.

This new model of trust is examined within the context of an e-commerce framework. The multiagent system is comprised of buyers and sellers that wish to conduct business. Sellers can engage in untrustworthy business behavior at the buyer's expense. It is the job of the model to decide whether a selling agent is trustworthy enough to engage in business. The trust model is analyzed with respect to stability, scalability, accuracy in attaining e-commerce objectives, and general effectiveness in discouraging untrustworthy behavior.

Based on the experiments, the model appears to be scalable dependent upon the agent population of buyers and sellers. It achieves its primary objective of discouraging untrustworthy behavior as measured through the acceleration of Gross Domestic Product growth over time. Within the simulator, a high degree of random outcomes is possible. Stability is used to examine the predictability of the model (on average) given a fixed set of given data about the simulations. Based on the simulations, the model appears to be quite stable.

1 Introduction

It is our view that the ability to reason with and about trust of various forms can be beneficial to the autonomous artificial agent. Aside from the trust internal to an agent society, another form of trust exists between the agent and the end user. The end user expects the Multiagent System (MAS) to fulfill their needs, and,

in the event that agents usefully and successfully perform a service, the user's trust in the society will rise. In both types of trust, our definition of trust related to reputation applies [1]. A foundation of trust between the user of the system and their agent counterparts will be established. The lack of this establishment will naturally lead to an undermining of the goals of an agent society that relies on user participation.

Given a large number of options, users will not use a MAS that they do not trust. Hence, a dysfunctional MAS may be established. Clearly, a functional system is needed to create the necessary trust and confidence of the user in the MAS. At the same time, both levels of trust are dependent on each other for the success of the system.

We propose that an explicit model of trust must be integrated into a MAS in order to allow for profitable dependency amongst agents.

1.1 Model Overview

Currently, many researchers are focusing their efforts on formalizing trust within MAS's based on an e-commerce framework. A formalization of trust will lead to models that accurately reflect the same characteristics and behaviors of trust that humans exhibit within human societies. Through this reflection, both levels of previously mentioned trust can be established.

There are as many proposed models of trust as there are people attempting to model it within the computer science discipline [16, 4, 1]. This is due in part to the fact that trust is both a social and personal phenomena [19]. Trust has the characteristics of being both objective and subjective. Most organizational scientists define trust as the willingness of a party to be vulnerable to the actions of another party based on the expectation that the other party will perform a particular action important to the trustor, irrespective of the ability to monitor or control that party. In contrast, psychologists tend to define trust as differences in personality. Clearly, trust-related disciplines (personality theory, behavioral decision theory, social psychology, sociological theory, and artificial intelligence) have varying perspectives on definition [27, 16].

Fortunately, there are many converging themes between these disciplines with respect to properties of trust. Most notably, trust is perceived as "good" while distrust is perceived as "bad" [27]. As well, social context is less emphasized in defining trust. Lastly, trust appears to be balanced and multidimensional in nature.

This work aims to establish a new model of trust for a MAS with a primary goal of clarification of the concept of trust. Based on previous definitions of trust, it is proposed that many separate concepts directly linked to trust such as reputation have been incorrectly included within trust definitions. This inclusion without labelling has led to the prevalent problem of large, unclear, and conflicting definitions of trust. It is proposed that trust is composed of several different phenomena that can be separately formalized and brought together to represent trust.

The new model proposes that trust is a combination of self-esteem, reputation, and familiarity within a MAS context. These concepts are linked to the idea of fulfillment. The model proposes that an agent's trust is based on fulfillment of roles, goals, and ideals of other agents. No known models have introduced the idea of self-esteem as it relates to agent trust.

The following highlights the properties of trust that are reflected in the proposed model:

situation specific	agent specific
objective	subjective
non-symmetric	multidimensional

Our previous research has already formalized reputation within the confines of an information sharing MAS [9]. It is desirable to lift the formalization to a higher layer of abstraction in order to apply the concept of reputation to all types of MAS's. The formalization of an agent's self-esteem will be defined at the same level of abstraction. The work provides examples of the use of the model in e-commerce and information-sharing agent societies.

The newest concept introduced to the theories surrounding trust involves the self-esteem. We believe that the quality of the mechanisms governing trust within an agent can be objectively measured through a sampling of the self-esteem of the agent as the agent judges the trustworthiness of other agents and observes the results. Based on the literature review, no objective measures of the quality of the mechanisms governing trust have been established.

The self-esteem of an agent indicates the degree to which an agent believes it has made a "good" judgement call. A "good" judgement call is based on the perception that an agent achieved some goal based on the actions of another agent.

Due to the fact that this approach is new, the work has a large theoretical component. However, a science/engineering approach is presented. A prototype is constructed to simulate trust and demonstrate proof-of-concept. The end result of the first component is algorithms that are used to describe the functioning of trust along with simulation results of a proposed game to demonstrate the concepts.

2 Philosophical Dispositions

2.1 Trust

Rotter defines trust as "a generalized expectancy held by an individual or group that the word, promise, verbal, or written statement of another individual or group can be relied on [4]". This is very much in line with the classical definitions of trust found in most common dictionaries like Webster's Revised Unabridged Dictionary.

Social psychologists tend to emphasize the concept of vulnerability when defining trust. As an example, Mayer et. al define trust as, "the mutual confidence that no party to an exchange will exploit another's vulnerabilities [21]."

Luhmann argues that trust is a social relationship subject to its own set of rules [20]. It occurs within a framework of social interaction and personality.

Trust can be better understood through an examination of its need within a MAS with Luhmann's perspective in mind.

2.2 Reputation

In order to understand the concept of reputation, it is useful to examine the definition provided by sociologists. Sociologists have studied the concept of reputation as it relates to human societies.

Within sociology, the concept of reputation lies within the realm of Sociology of Identity. In turn, Sociology of Identity is connected to the outside world through Symbolic Interactionism theory. This provides the logical connection between multiagent studies within the Artificial Intelligence (AI) community and the studies of identity and reputation within Sociology.

Sociologists see reputation as held by individuals, organizations, and objects. It is ascribed by society towards the individual and cannot be generated by the individual being judged. Reputation is viewed as both personal and commercial.

Personal reputation is the general estimation held by the public about the individual in question. This estimation is based on the way in which identity is managed by the individual and presented towards society. Through presentation of the self, society constructs a reputation by looking at the objective behavior of the individual [14].

Commercial reputation is very similar to personal reputation in that it focuses on the estimation held by the public about the business in question. However, companies utilize public relation components to influence how society perceives them. Hence, companies have more power than the individual in establishing and marketing their own reputation. This helps to explain the unequal distribution of social power between companies and individuals.

Erving Goffman makes the useful analogy of identity management to the dramatic world of theater [14]. Individuals become actors on a stage negotiating their roles to society. On the front stage, individuals are polite with decorum. The back stage is where the individual exists as themselves. Front stage and back stage attain a certain fluidity and reputation is created through the management of both stages.

The audience has certain expectations or roles for the actors. Reputation is constructed based on the audience's belief that the actors have fully satisfied their roles. If an actor cannot act, they have failed as the role of actor and will develop a negative reputation amongst the audience members. In the same way, individuals and businesses have certain roles within society. If society objectively judges that they have met their roles, they are rewarded with a positive reputation.

Positive reputation leads to confidence/trust in the individual along with a higher level of social status and power [7]. Such individuals become sought out within society. Negative reputation leads to a loss of esteem held in society along

with social status and power [30]. Naturally, it is in an individual's best interest to maximize their positive reputation through identity management.

Within human societies, both reputations are so important that rules are established to protect reputation [13]. Defamation laws prevent an individual from destroying another individual's reputation through false accusations with third-parties [12]. This suggests that within a MAS, the issue of defamation must be addressed in the formalization of reputation. Separate work from Elijah Bitting addresses this issue [5].

2.3 Roles

In a nutshell, self esteem is the reputation an individual acquires with themselves [6]. This has very powerful ramifications on the understanding of self esteem as it requires a good understanding of the concept of reputation.

As discussed earlier, reputation is an expression of role fulfillment. The selection of roles and ways of measuring them constitutes the challenge of defining reputation. Within psychology, self esteem expresses the perception of how an individual measures up to satisfying their personally defined roles [6]. Individuals can suffer from low self esteem in one of two ways. They can either select inappropriate roles to fulfill or they can fail to perceive the degree to which they are fulfilling those roles.

As an example, consider the case of an individual that thinks they should be a flawless ice skater. If the individual is not a professional ice skater, they have chosen a wrong role to include within their self esteem. They will inevitably fail at living up to their expectations because they are not a professional. Self esteem will suffer as a result. This exemplifies inappropriate role selection.

If the same individual is a mother, she may perceive herself as an inadequate mother despite external evidence contrary to the notion. In such a case, self esteem suffers because the perception of role fulfillment is incorrect. Hence, the challenge in maintaining self esteem is to select adequate roles and ways of perceiving the degree of satisfaction.

Role is a key concept in role theory. Sociologists rely on role theory to provide a definition of this concept through two conceptually different accounts [15].

Medical sociology contributes to role theory through a structural definition of role. According to Ralph Linton, a role is an institutionalized cluster of obligations and rights that are normalized [8]. Linton offers the example of the sick role. A sick person is ascribed the sick role by society. Through ascription, a person with a sick role is allowed to take time off work within a certain given time frame. Employers must give this right while employees fulfill their obligations of recovery.

A similar definition that contributes to role theory is provided by social psychologists through a similar structural account. Social psychologists argue that a role is a bundle of rights and duties associated with a particular status within society. Although social psychologists do not formalize the concept of a status within society, they argue that an individual is comprised of many statuses within a society (status set). Each status, in turn, corresponds to a given role.

A typical example would be a doctor. A doctor has a recognized status within society. At the same time, a doctor has a given set of duties and rights with patients that must be honored in order to maintain the doctor role within society. Furthermore, a doctor may also be a parent and a teacher. An individual can subscribe to many statuses and corresponding roles.

The distinction between these two definitions of roles lies in the notion of role ascription. Social psychologists argue that roles can only be ascribed to individuals of a society [7] while social anthropologists argue that this limit on ascription is artificial and does not help them in explaining certain phenomena [29]. Fortunately, there is much consensus on the overall definition of the concept of a role.

Based on the author's previous work in formalizing trust [9], the social psychologist's perspective of roles will be more closely adhered to within this model of trust.

2.4 Values

The concept of a value is inherently vague and has meant many different things to different scholars. Sociologists and social psychologists use surveys to try and extract values common to a given population. Unfortunately, much research has been done to suggest that it is not useful to distinguish between personal values and opinions [23]. In essence, these surveys measure attitudes and opinions just as effectively as they measure values of the subjects involved. This only adds to the confusion over a precise definition of value.

On top of the definition, there is much uncertainty about the existence of specific values. Many questions arise about the specific values.

- How many values exist within a society?
- How do values relate to each other?
- It is more useful to look at a set of 500 versus 10,000 values?
- Do particular values form a fundamental set?
- If so, how do values form these sets?

Clearly, the nature of values is elusive as these are just a few of the many questions surrounding values. Fortunately, many disciplines have a limited but clearly defined concept of value. From these disciplines, there exist a clearly defined dichotomy of definition.

Traditional economic value theory contends that a value is the perceived exchange or market price of a given commodity. Economists distinguish between the value of a commodity and the value of an individual or society through the use of welfare economics. Within welfare economics, the value of an individual or society is realized through appropriate resource allocation. Essentially, individuals or societies are reduced to preferences [23]. This perspective of value is not useful to connecting value to trust.

Most other disciplines define value with an abstract, intangible meaning that is nonmeasurable. Sociologists construct value as a cognitive representation of

human need [26]. Value becomes defined as a belief with the intent of leading to the greatest satisfaction of the object. Here, objects are defined as individuals or groups.

Rokeach offers the following definition of value:

A value is an enduring belief that a specific mode of conduct or end-state of existence is personally or socially preferable to an opposite or converse mode of conduct or end-state of existence. A value system is an enduring organization of beliefs concerning desirable modes of conduct or end-states of existence along a continuum of relative importance [3].

Rokeach goes on to differentiate between two different types of values. Terminal values are those that relate to end states of existence. In other words, these values focus on the preference for a state of being after performing some given action. Examples of such values include the following: happiness, freedom, and economic wealth.

Instrumental values are those that are related to modes of conduct. These values focus on the behaviors leading to a desirable outcome. Examples of instrumental values include: honesty, benevolence, and commitment to goals.

Within terminal values, a second layer of abstraction allows for personal or social values. These values may be self-centered or society-centered. Values such as peace of mind or salvation are deemed self-centered. Values like brotherhood or world peace would be deemed as society-centered.

Lastly, Rokeach divides instrumental values into moral values or competence values. Moral values are seen as instrumental values centered about the self that arouse feelings of guilt when violated. The violation of competence values leads to feelings of inadequacy rather than guilt. Hence, moral values tend to be more intra-personal while competence values tend to be more interpersonal in nature.

Some anthropologists argue that values are a set of fundamental beliefs that serve to answer questions relating to the meaning of life and the nature of an individual's environment [18]. The answers to these questions serve to structure all subsequent decision-making and consequent behavior.

Value theory is branch of philosophy that focusses on the formalization of the concept of value. Value theorist Nicholas Rescher argues the following definition of value:

A value represents a slogan capable of providing for the rationalization of action by encapsulating a positive attitude toward a purportedly beneficial state of affairs. [25].

In other words, values act as a central means of rationalizing action within the human mind. Given a goal, values dictate the way in which the goal will be accomplished. This definition agrees with the sociological and anthropological definitions provided above.

The modelling of trust can be abstracted to the modelling of a decision-making process. In this case, the decision involves the ability to trust another

agent. Hence, one can make the argument that value theory is fully applicable to modelling trust.

Throughout this work, a combination of the sociological definition of value along with value theory will be used to model values within trust.

2.5 Self-Esteem

Based on the assumed definition of reputation, self-esteem is defined as the reputation an agent has with itself. This definition of self-esteem is widely accepted and does not require further elaboration.

2.6 Familiarity

Relational psychologists argue that familiarity between two agents is a result of similarity in underlying value-systems shared by the two individuals. For our work, this same disposition is held.

3 Generalized Model

Figure 1 denotes a concept graph that illustrates the set of dependencies amongst the concepts previously discussed.




Fig. 1. Concept Dependency Map

As can be observed, trust can be defined as being dependent on an agent's reputation. If an agent has a perceived higher reputation, another agent will tend to trust them more with respect to the context of the reputation. The careful labelling of reputation as being perceived is intentional. Perception is innately subjective in nature. To perceive reputation is to perceive that its roles are satisfied. Hence, reputation acquires a subjective aspect through its perception. At the same time, social reputation becomes objective through consensus of definition and observation.

Reputation, in turn, is dependent on the roles that are used to define it. In reality, each individual is allowed to have a different set of roles and expectations for another person. Reputation of the other person occurs on a personal or social basis. Personal reputation is based on the enforcement of role ascription of one individual onto another while social reputation is ascribed by society based on a set of commonly agreed upon roles.

Roles act as a manifestation of values. Values are a subset of the beliefs of an individual or society that express the preference for one state of events over another in the future. The definition and selection of value is innately vague. As seen later, other disciplines will help in defining and selecting appropriate values for the model.

Although trust can be thought of as conceptually dependent on reputation, it is multidimensional in nature. It can depend on many other things. Trust is already defined as being indirectly dependent on values through reputation; it is believed that trust can be directly dependent on values as well.

The direct dependence of trust on values can be argued through Value theory [25]. As stated earlier, trust can be further abstracted as a decision making process. Value theory concerns itself with the impact of values on a given decision making process. In essence, values act as a rationalization mechanism for making decisions. This implies that trust can be chosen for a multitude of different reasons that directly depend on an underlying value system.

The highly subjective rationalizations that can occur behind trust are explainable if values are taken into account. Due to the hierarchical nature of values, the task of deciding to trust becomes simplified if it is based on an individual's values.

As an example of subjective rationalization, consider the case of a male automobile mechanic that may be highly skilled in fixing BMW's. This same mechanic may be generally perceived as chauvinistic towards women. A female customer may be more inclined to not trust the mechanic to fix her car if she is aware of this strong difference in values. Meanwhile, the male customers may value the mechanic's skills more than his beliefs towards women. Although this is an arbitrary example, much evidence supports the notion that similarity in value systems does have a significant impact on the process of trust between two individuals [26, 2, 20].

One can consider the choice of relying on reputation over similarity towards trust as yet another expression of individual values. Social reputation reflects what it is commonly agreed upon (collectivism) while similarity reflects indi-

viduality. Due to the hierarchical nature of values, it is not possible to value both equally. Hence, individuals that wish to maintain group harmony may focus more on social reputation; those that focus on individuality may be more inclined to trust based on individual experiences or perceptions [10].

In the same sense, these arguments can be connected to a MAS through agents and their respective reputations. Users of agents possess a hierarchy of values that are reflected through the definition of an agent. The relative importance of different values is highly subjective in nature. Hence, similarity of values can be measured through the Hamming distance of agent value hierarchies. This allows an agent to rationalize trust in the same way as the user; otherwise, we conjecture that users will invalidate the agent's trusting mechanism.

The issue of the formalization of reputation has already been addressed in previous work of reputation formalization [9]. However, the introduction of values into the trust model allows for a more systematic approach to defining the roles and subsequent reputation. Previously, the roles were defined using a rather ad-hoc approach.

Lastly, the concept of self-esteem is introduced to the trust model to allow for value adaptation. As an agent satisfies roles that reflect its value system through actions, an agent has the opportunity to assess the validity of its value hierarchy. Does its trusting mechanism result in desirable end states that the user wishes? If so, the relative importance of values is correct. Otherwise, values must be changed in importance so that trust is exercised correctly as the user sees fit.

Self esteem acts as an assessment of the trustworthiness of an agent in its own trusting mechanism. Without this feedback mechanism, trust model invalidation may result from the agent's owner. As Luhmann states,

Trust can only be maintained if it finds a form which allows it to live with such suspicion and be immune to it [20].

Below, each component of the hierarchy will be fleshed out in detail within the contexts of our e-commerce MAS.

Once values have been defined, they can be manifested through actions or objects. Actions take on a predictable set of behavioral patterns referred to as roles. This section presents the roles that represent the connection between the highly abstract values and trust within society and the individual.

Figure 2 denotes the proposed roles as nodes in a graph. Each level of depth represents a layer of abstraction in role definition. The leaves of the nodes represent the values that are manifested. A discussion of the proposed hierarchy is presented below.

3.1 Role Patterns

Conceptually, the agent can be seen as satisfying many different roles. The choice of which roles to highlight introduces an element of subjectivity into the proposed

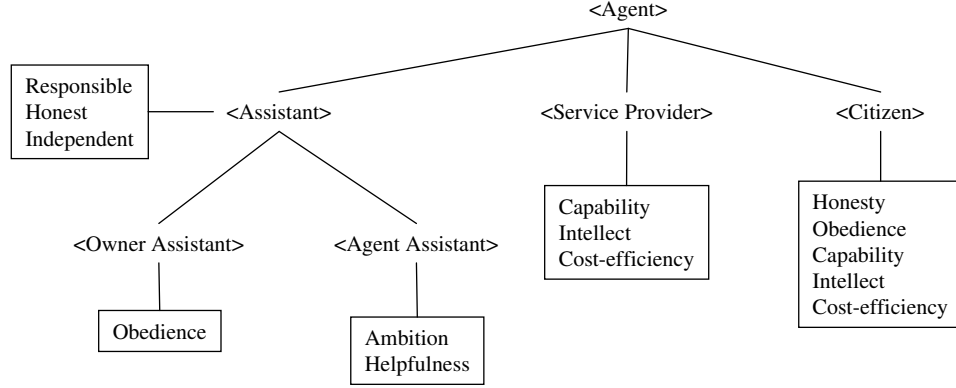


Fig. 2. Value Manifestation Tree

model. In this case, the chosen roles were based on the agent definitions provided in Section 2.

There are a few safe assumptions implied within this model. First, it is assumed that an agent can interact or communicate with fellow agents or an agent's owner. Additionally, this same implication implies that an agent has an owner. Second, it is assumed that the agent is a member of a greater society or MAS. Lastly, it is assumed that the agent provides some meaningful service for others.

3.1.1 The Assistant

An agent can be seen as an assistant, a service provider, or a citizen. The values of responsibility, honesty, and independence are embedded directly within the role of an assistant. These values imply the following desirable qualities of any assistant: dependability, reliability, honest, self-reliant, and self-sufficient.

Conceptually, an assistant can be an assistant to its owner (the user) or another agent. If an agent is an assistant to another agent, the values of ambition and benevolence are useful to have in addition to those of any assistant. They are good indicators of the willingness of an agent to help another agent. An agent that is seen as an assistant to an owner must value obedience on top of the other qualities of an assistant. An agent must maintain high priority on attempting to achieve the goals issued by its owner relative to satisfying other agents' requests.

It is worth noting that the values of independence and obedience do not conflict within the owner assistant role. Independence is defined as self-reliance from other agents while obedience focuses on adherence to goals set by the owner. There does exist a conflict between being a responsible assistant to the owner versus a fellow agent. If an agent fails to be obedient to the owner, it becomes more reliable as an agent assistant. If an agent becomes more obedient to the owner, it abandons goals agreed upon by fellow agents. In doing so, it

becomes less responsible as an agent assistant. Unfortunately, an agent cannot be everything to everyone.

3.1.2 The Provider

Next, an agent can be seen as a service provider. A service provider must value capability and intellect. When both values are combined together, the end product is a high quality service. A service can be measured as being high quality if it meets or exceeds the needs of the entity requesting the service. A low quality service reflects a deficiency in either capability, intelligence, or both. Separately, these values are much more difficult to measure through actions because of their innate vagueness. Hence, they were grouped together in a ‘value basket’ [11].

The previous research conducted in the formalization of reputation [9] could be considered as a focus on the agent as a service provider.

Normally, one would expect an assistant to be conceived as a service provider. Within the tree, the assistant would not be on the same level of abstraction as the service provider. In this case, the division between assistance and provider has been made with the assumption that the service will be performed in the same manner independent of who is being served. In other words, an assistant is conceptually tied to who is being served while the quality of service is not. Hence, the assistant can be separated from the provider under this assumption. Otherwise, the service provider node would fall one level below the assistant node.

3.1.3 The Citizen

Lastly, an agent is assumed to be a member of a MAS that contains its own rules and social values. These values may or may not conflict with the individual values of an agent. Society wishes to promote those values that preserve itself.

Society values honesty in order to facilitate trust. This need for trust within society has already been established in Section 2. Society values self-control in order to facilitate resource preservation. Resource preservation is clearly needed as resource depletion leads to degradation in agent functionality. The term ‘resource’ is deliberately left unspecified because it may mean different things in different MAS’s. As an example, a resource within an e-commerce society may be an economy while a resource within an information-sharing society may be information.

Lastly, society values the capabilities of its agent. It is conjectured that agents that are highly capable and honest will be utilized more often than not. This utilization directly translates into increased utilization of the society by agent owners. In essence, the society should encourage all of the values previously established within the agent as a service provider.

3.2 Role Fulfillment Measurement

As discussed in previous research on the formalization of reputation [9], each role must be defined and a way of measuring its degree of role fulfillment must

be established. When calculating reputation, one must be able to ascertain an individual’s ability to perform an ascribed set of roles.

This measurement introduces an element of subjectivity and objectivity into the process of reputation attainment. The measurement is subjective in the sense that it is based on a belief that a particular set of actions will reflect role fulfillment. Furthermore, the observation of the actions may be based on a subjective perception. The measurement is objective in the sense that its calculation is based on an algorithm. The previous research’s algorithms were based on counting techniques of observed patterns of behavior.

In the previous research, role fulfillment measurement was generally achieved through observation of a given behavior. This behavior becomes represented as an instance of role satisfaction that decays over time. The total degree of role fulfillment becomes a summation of each of these instances over a given window of time. Most roles had a unique way of calculating the impact of one instance of satisfaction on the overall role satisfaction.

From a pragmatic perspective, the task of trying to measure the degree of role fulfillment of these roles seems quite daunting. The chosen actions to observe must reflect the underlying values. Furthermore, the choice of how to calculate the degree of role fulfillment must be agreed upon.

Below, a general overview of how measurement takes place within each role is presented in the same way as with the concept map. \mathcal{A} denotes the role of the assistant. \mathcal{P} denotes the role of the provider. Lastly, \mathcal{C} denotes the role of the citizen.

3.2.1 The Assistant \mathcal{A}

Conceptually, it is believed that all assistants minimally reflect the values of responsibility, honesty, and independence. These are desirable qualities of both an owner assistant and an agent assistant. An itemized list of the values and their measurable actions is presented below. \mathcal{A} represents the degree to which this role has been satisfied with respect to these values.

- **Responsibility, v_1 :** Responsibility really reflects the dependability of an agent as an assistant. Externally, one can observe the number of times an agent agreed to achieve a goal and abandoned it. A responsible assistant will not agree to take on a goal unless it is confident in its ability to accomplish that goal. Hence, responsibility can be measured through observation of this behavioral pattern.

Responsibility is formalized through observation of the total number of attempted service contracts c_a to the total agreed upon service contracts c_t .

$$v_1 = \frac{c_a}{c_t} \quad 0 < v_1 \leq 1 \quad (1)$$

The source of these contracts is irrelevant as the quality of responsibility should not depend upon the entity being served. Furthermore, the magnitude of the load is irrelevant to the measurement.

Within an e-commerce framework, c_a denotes the number of attempted services performed by a bidding or selling agent. c_t denotes the number of services agreed upon by that same agent. Initially, v_1 must be set to zero until c_t is greater than zero.

- **Honesty, v_2 :** This work is currently being developed in conjunction with work related to defamation [5]. An agent that is honest will not defame a fellow agent. Through [5], a mechanism for the detection and handling of defamation is being studied. One component of honesty, α_1 , can be measured through the number of defamations that have been established against an agent.

$$\alpha_1 = \left(\frac{d_p}{d_t} \right)_i \quad (2)$$

The ratio of the number of proven defamations d_p from an agent i to the total number of accusations of deception d_t acts as a mechanism of honesty measurement. Defamation formalization leads to a useful and practical measurement of honesty. An agent that is honest will not attempt to defame users or their agents.

Another component of honesty, α_2 , can also be related to the quality of work. An agent that performs honestly does not sacrifice accuracy for convenience. It does not misrepresent its services. Hence, users must be allowed to subjectively express disapproval with an agent's services. The ratio of the number of dishonest service reports s_d to total service reports s_t for agent i represents a useful measure of this type of honesty.

$$\alpha_2 = \left(\frac{s_d}{s_t} \right)_i \quad (3)$$

There is always the risk of dishonest expressions with the intent of defamation. In such cases, defamation must be dealt with and handled appropriately. Separate work performed in parallel by Bitting focuses on the formalization of defamation [5]. As such, this work will not address the issues surrounding defamation. Instead, the knowledge of defamation will be treated as a black box within a repository accessible by the buying and selling agents.

$$v_2 = w_1\alpha_1 + w_2\alpha_2$$

$$\begin{aligned} w_1 &= \begin{cases} \frac{\alpha_1}{\alpha_1 + \alpha_2} & \text{if } \alpha_1 \geq \alpha_2 \\ 1 - w_2 & \text{otherwise} \end{cases} \\ w_2 &= \begin{cases} \frac{\alpha_2}{\alpha_1 + \alpha_2} & \text{if } \alpha_2 \geq \alpha_1 \\ 1 - w_1 & \text{otherwise} \end{cases} \end{aligned} \quad (4)$$

Each component α_i is weighted with w_i and the dot product of components and weights produces a scalar value v_2 . The weighting is a function of magnitude as means of balancing defamation against subjective expression.

- **Independence, v_3 :** Independence implies self-reliance. Wherever possible, an independent agent will rely on itself to achieve its goals. A dependent agent will unnecessarily rely on services offered by others that it could perform itself. It is possible to measure independence through the examination of requests for services by one agent to another. A dependent agent will request services that it can perform itself. Formally, independence is measured through the following mechanism:

$$v_3 = 1 - \left(\frac{r_u}{r_t} \right)_i \quad (5)$$

where r_u represents the total requests issued by agent i that are deemed unnecessary, and r_t denotes the total services requested by that agent. This formalism implies that the system must have a predefined knowledge of the service capabilities of each type of agent. We conjecture that this is a valid assumption within the context of our e-commerce MAS.

This value is an indicator of other things besides independence. An agent that is overloaded in work may decide to offload its work onto other agents that are idle. Hence, there may be conflicts between the values of independence, capability, and intellect. This is acceptable due to the nature of values as being potentially conflicting.

There exists two distinct assistant types. An owner assistant is an agent that directly helps the owner while an agent assistant aides other agents. It is possible that a single agent can satisfy both roles. Below, a continuation of the measurement process for the values of an owner assistant is outlined.

- **Obedient, v_4 :** Ideally, services requested by the owner of an agent would be performed as quickly as possible. This implies that an owner's requests have priority over other agent's requests. Obedience to the owner can be measured through an examination of consistency in achieving the owner's goals over those already within an agent's waiting queue of tasks. The definition of obedience implies a time and order dependency. Agents must satisfy their owner's goals as quickly as possible. Furthermore, agents must execute the goals in an order issued by the owner. Formalization of obedience is as follows:

$$v_4 = \frac{c}{o} \prod_{i=1}^n (1 - t_i). \quad (6)$$

c denotes the total goals achieved for the agent's owner while o denotes the total goals issued by the owner. Ideally, this ratio will equal one as it denotes total achievement. Next, the product denotes the penalty incurred for not satisfying the owner's n requests within a given window in strict order. t_i denotes the delay due to satisfaction of other agent requests for a service before the owner's i th request. Hence, t_i has a cumulative effect. As other agent's requests are satisfied above an owner's service request, t_i approaches 1 and the measure of obedience quickly approaches zero.




Fig. 3. Process Queue Example

This value does not imply that services performed for other agents must be executed in strict order. This value is impartial to others' requests.

Figure 3 demonstrates an agent's current list of goals to achieve. Process p_i represents a request from a fellow agent or the owner. Within this definition of obedience, one would expect processes p_1 and p_4 to be executed first so as to minimize delays for the owner. Processes p_2 and p_3 would move down while p_4 would move up. In this case, suppose things get executed in the order already in the queue. t_1 is equal to zero because of an immediate execution while t_2 is equal to the product of the time delays incurred from p_2 and p_3 . Although processes p_1 and p_4 did complete, a penalty is incurred due to the accumulative delay of processes p_2 and p_3 .

Within the context of our e-commerce MAS, agent requests may take the form of price checks, recommendations, or transaction verifications. Owner requests take the form of buying or selling of a tangible good.

Lastly, the practical way of measuring the distinct values associated with an agent assistant are outlined below.

- **Helpfulness / Benevolence, v_5 :** Agents should be willing to do things for other agents without an immediate pay off. The benefits to society can be enormous [22]. The labelling of an act as benevolent or not is vague as it is subject to speculation. Despite the labelling, the same outcome results.

Hence, one can measure helpfulness or benevolence through the volume of requests to denials of service made by an agent.

Formalization of this value comes through an examination of the services performed to satisfy agent requests n_a out of the total services n_t within a given β frames. Essentially, v_5 is the average of this ratio over the given frames:

$$v_5 = \frac{1}{\beta} \sum_{i=1}^{\beta} \frac{n_a}{n_t} \quad (7)$$

The size of the frame must be large enough to allow sufficient time to carry out n_a requests. For subsequent experiments, each frame is arbitrarily set at one day.

3.2.2 The Provider \mathcal{P}

As a provider, an agent has an obligation to the MAS to provide as high a quality of service as possible. Other agents and the owner may need to rely on them at some point in the future. Failure to provide a service will result in a loss of confidence in the MAS and the agent. Below, measurement techniques are outlined for evaluating the role of a service provider. \mathcal{P} represents the degree of satisfaction of this role.

- **Capability/Knowledgability, v_6 :** As discussed before, the value basket of capability and knowledgability is created due to the difficulty in distinguishing between one and the other in external behavior. Together, these values express themselves through the ability of an agent to perform requested services that should normally be expected of an agent. Furthermore, agents and owners will be satisfied with the quality of a highly capable agent. Measurement occurs through the observation of expected performance capability versus volume of performance deliverability. Furthermore, issues of quality of service may be connected to values of honesty. If an agent is dishonest in their services, they fail both capability and honesty. As an assistant, they are punished. As a service provider, they are further punished. In order to formalize and measure this value within an agent, a chosen action must be externally observed and its frequency measured. One considers the quality of a service as a manifestation of both capability and intellect. The two values cannot be easily distinguished. Hence, a value basket allows one to combine both values and measure their combined effect. In this case, the difference in expectations between a service requestor and the service provider acts as a practical measure of the satisfaction. An agent that meets or exceeds the expectations of the requester has provided an excellent service. Otherwise, the requester is disappointed. Let E_p denote the expectations of a service provider for service X at time t . Let E_r denote the expectations of a service requester for the same service X

at time t . The satisfiability of the service, Γ , at time t between the provider and requester is defined as follows:

$$\Gamma(t) = E_p(t) - E_r(t) + \epsilon \quad (8)$$

where ϵ represents the uncertainty factor within the expectations. If $\Gamma(t) = 0$, satisfaction has been met and the provider is deemed capable of performing the service X at time t . If $\Gamma(t)$ is greater than 0, the service provider has exceeded the expectations of the requester and is deemed more capable and knowledgeable within that domain. Lastly, $\Gamma(t)$ less than 0 denotes an unsatisfied requester with a perception of incapability.




Fig. 4. Relationship between ϵ and $\Gamma(t)$

Figure 4 denotes the relationship between $\Gamma(t)$ and ϵ . In this graph, ϵ represents a partial membership function of satisfiability. If $\Gamma(t)$ falls within $\pm\epsilon$, the service has been satisfied. Any value greater than this range guarantees over satisfiability while any value less than this range guarantees unsatisfiability. Essentially, $\Gamma(t)$ becomes fuzzified.

The formalization of the expectations E_p and E_r is domain-specific as expectations are dependent upon the service requested. An intangible service would be more difficult to formalize than a tangible service as intangible services cannot be externally observed by definition. Hence, this work focuses on tangible services with an avenue for future work surrounding the issues of intangible services. E_p and E_r are dependent on priorities, capabilities of the provider, and knowledge. Naturally, E_p and E_r must be of the same scale in order for this to work.

Within an e-commerce framework, E_p and E_r denote the expected prices associated with the buying and selling of tangible goods. A selling agent is observed to sell a tangible good at price E_p while the agent owner wishes to sell for no less than price E_r . If the selling agent can sell for significantly higher than E_r , the agent has manifested a higher degree of capability and knowledge of the market. The value of E_r is based on the average market value of the good. E_p results from an internal strategy specific to the agent. Combined, these differences in expectation and satisfaction are averaged over n samples to produce a scalar value:

$$v_6 = \frac{1}{n} \sum_{t=1}^n \Gamma(t) \quad (9)$$

3.2.3 The Citizen \mathcal{C}

Lastly, an agent is a member of a greater society. It must obey the values of society or be punished as it will be directly harming society's existence. The section presents the ways of measuring the underlying values of a citizen. \mathcal{C} represents the degree of satisfaction of this role given the values that it implies. It will be explained in detail within the reputation section.

- **Cost-Efficiency, v_7 :** Cost-efficiency enforces effective and efficient resource utilization. In the physical sense, an agent consumes memory, disk space, and processor time. Examinations of these patterns of use could reveal unacceptable resource allocation as deemed fit by the MAS designers. Subsequent penalties would be enacted to encourage less consumption. Formally, cost-efficiency manifests itself through the maximization of the utility function u_i^t of given resources R for services ς offered by the agent i at time t . u essentially measures the effects of resource control.

$$u_i^t = \text{value}(\varsigma) - \text{value}(R) \quad (10)$$

Cost-efficiency happens when u_i^t is greater than or equal to 0. Due to the nature of this definition, services and resources are dependent on the nature of the MAS. Hence, formalization of ς and R cannot be universal.

Within an e-commerce framework, definitions extend to define the value of ς as the difference in the cost of the tangible goods bought or sold relative to an established average cost. Resource R is defined as the cost of a computational resource like processor, memory, or disk space. The formalization of R must be built upon common consensus.

Combined together, this value is a result of the running average of M transactions:

$$v_7 = \frac{1}{M} \sum_{t=1}^M u_i^t \quad (11)$$

- **Honesty, v_2 :** Honesty is derived from the same measurements carried out when considering the agent as an assistant (See Section 3.2.1 and Equation 2).
- **Capability/Knowledgability, v_6 :** Capability and knowledgability are measured in the same way as when considering the agent as a service provider. A MAS wishes to attain highly capable agents as it directly translates into increased utilization. (See Section 3.2.2 and Equation 8)

3.3 Reputation

In the previous subsection, the expected roles of the users were formally expressed and models were established to express the degree of satisfaction of each of these independent roles. Note that there are other ways to define, measure and quantify these roles, some of which could be probabilistic in nature. One should realize that the definitions and measurement schemes presented in this work are an extension to [9].

In order to determine a user's reputation, one must examine the degree of fulfilment of each of these roles and combine them to produce a scalar value reflecting the final summation of degree of satisfaction of these roles. Let

$$\mathbf{R} = \{r_1 = \mathcal{A}, r_2 = \mathcal{P}, r_3 = \mathcal{C}\}$$

and

$$\mathbf{W} = \{w_i \mid i = 1 \cdots 3\}$$

represent the corresponding weight vector. We define the user's reputation, \mathcal{R}_u , as

$$\mathcal{R}_{(u)} = \sum_{i=1}^3 w_i r_i = \mathbf{R} \cdot \mathbf{W} \quad (12)$$

The above role vector, \mathbf{R} , expresses the degree to which each of the previously mentioned role patterns has been satisfied independently of the others. The normalized weight vector, \mathbf{W} , expresses the degree to which each role's satisfiability must be weighted in calculating a final reputation. Each vector component r_i is defined iteratively the same way. \mathcal{A} , \mathcal{P} , and \mathcal{C} are each calculated by taking a similar dot product of each of their values and corresponding weights and producing a scalar value to represent their particular role satisfaction.

The distribution of weights on both levels is a reflection of the importance of values within an individual agent or MAS. A MAS may place more emphasis on the value of responsibility than service performance. In such a case, an agent will weigh the role of an assistant more so than the role of provider. Once again, these values are entirely dependent on the goals of the MAS and end-user.

In [9], it was believed that it was impossible to universalize the calculation of reputation as the weighting function is dependent on values. In this work, the values have been universalized to all societies through the model above. While the importance or weighting function of values cannot be applied uniformly to all agents and societies, the values do exist within all societies. Hence, reputation is made more universal than before.

Within the formalism, the dot product of the role vectors will result in a scalar value that expresses the magnitude of the reputation. This overall reputation must then be normalized between [-1,1] using a hyperbolic tangent function.

It is very important to capture the semantics of the roles and determine the significance of each role in relation with other roles. The weights can be determined in a number of ways. One can calculate the weight of a role by computing its statistical significance, \mathcal{S} , as follows,

$$\mathcal{S}_{r_i} = \frac{r_i - \text{mean}(\mathbf{R})}{\sigma^2(\mathbf{R})}$$

The weight distribution vector becomes the normalized representation of this new statistical significance vector. The roles which deviate most significantly from the measured mean become the biggest contributors to the calculation of reputation.

For the purposes of the experiments conducted later on, the weight distribution function relies on an initial setting by the agent’s owner. As discussed in Section 2, the agent must reflect the values of the user as accurately as possible in order to make valid trust decisions from the user’s perspective. Otherwise, incongruence in value systems between the agent and owner will lead to user invalidation of the trust model.

3.4 Trust

Trust can be exercised for many reasons through different mechanisms. Within the context of the previous work [9], trust is exercised based solely on reputation. An agent with an excellent reputation is guaranteed to be trusted by everyone. An agent with a recent dubious past is assured of social ostracism. While this is generally true within human society [24], other factors can come into play when exercising trust.

[9] fails to address important aspects of a relationship between two agents. Primarily, it does not address the trust that results from friendship. Within friendship, the decision to trust can be extended beyond the parameters of previous experience. In other words, friends can trust each other to do things that lie outside of an established reputation.

As an example, two individuals may be friends for several years. Both friends have a lot in common and work in the same field. They trust each other enough to borrow money from one another with the assumption of being paid back. Based on the previous trust model, such ‘leaps of faith’ are not possible if the friends do not have a credit history of each other. Each person would subscribe to the role of financial creditor and a history of transactions would establish the needed role fulfillment for trust to occur. It is not possible to explain this particular mechanism of trust using the previous model.

Relational psychologists argue that more often than not, opposites do not attract [17]. Friendships between individuals are established on similarity [10]. Specifically, similarity is expressed as closeness in beliefs and attitudes of the individual towards themselves and the external world. These personal beliefs and attitudes directly express the individual’s underlying value system. Consequently, value system similarity has been shown to be a reliable predictor of friendship between two individuals [10]. At the same time, most psychologists agree that trust is an essential feature of friendship [31]. Essentially, one cannot have friendship without an element of trust. The opposite has also been found

to be true. Trust between two individuals cannot exist without some level of friendship [31].

Based on this argument, it is possible to trust by examining the similarity in the underlying value systems. Although agents within this model share the same values, the hierarchy of importance of the values is personalized. Hence, trust between two agents can be exercised through a measurement of this similarity. Essentially, the Hamming distance between the value hierarchies acts as an expression of the degree of friendship and resulting trust. Two agents that have identical hierarchies will trust each other more so than agents with opposite hierarchies.

Formally, the Euclidean distance between two different agents' hierarchy vector V can act as a measure of similarity between the two agents with respect to values. After the distance is calculated, the resulting vector must be divided by the maximum distance vector to express membership degree. Alternatively, the inner dot product of the two vectors can be combined with the cosine of the angle between them to act as another measure of similarity.

This model introduces yet another tool for deciding upon trust of an individual. The choice of which tool to use in evaluating trust depends on the context. An agent that has previous knowledge (personal and social experiences) of the other agent may wish to rely on reputation. In the event of total ignorance, friendship assessment should be relied upon more often.

As a convincing example of this principle, consider the example mentioned above. Two friends may initially lend each other money based on their relationship status because they have complete ignorance of each other's lending history. Over time, they build a lending history and utilize reputation to lend each other money. If one friend is terrible at paying the other back, they may remain friends but do not loan each other money. The 'leap of faith' transitions to a rational process. Initially, they exercised trust based on friendship assessment. Once they had experiences, they exercised trust based on reputation.

3.5 Fuzzy Logic Integration

The choice of using numerical representation to measure values, reputation, and trust demands one to consider the nature and ramifications of numerical representation within the model. Each type of value is measured and a scalar value results. These crisp values are then linearly combined to form reputation with the assumption of familiarity. Hence, reputation becomes a crisp value along with the resulting trust.

Boundaries must be set with each of these crisp values to determine the degree to which each element is fulfilled. At the highest level, trust has undefined upper and lower limits. Within this range, an agent must be perceived as having a degree of trustworthiness. Several different methods can be applied to define these acceptable boundaries. An arbitrary constant can set an upper limit. Any trust measurement beyond this upper limit allows the agent to trust the other agent. Any agent falling below that threshold value is deemed untrustworthy for the service requested.

This technique is not very useful as it relies upon fixed membership. A given trust value falls into one of two sets: trustworthy, or untrustworthy. An infinitely small difference between two given trust values could result in different membership assignment. Naturally, this does not make sense as humans do not evaluate trust based on such black-and-white terms. This is due to the fact that this linguistic term can be seen as qualitative as well as quantitative. The example of economic poverty versus richness is yet another example of the weakness of using this method [32].

This trust model must reflect an element of uncertainty or vagueness in order to have human validation. In other words, membership partiality within trust must be allowed. The equation below illustrates a membership function μ that describes the classification scheme behind the value of trust t .

$$\mu = \begin{cases} 0 & \text{if } t \leq \alpha_1 \\ \alpha & \text{if } \alpha_1 < t < \alpha_2 \\ 1 & \text{if } t \geq \alpha_2 \end{cases} \quad (13)$$

$$\alpha = 1/(\alpha_2 - \alpha_1)$$

Membership function definition is dependent on the scalar values of α_1 and α_2 . By setting these values beforehand, a gradient in membership α is allowed for all potential values of trust. α represents a constant slope. There are several techniques possible for deciding on appropriate values. It is common to set the values based on expert knowledge of the related problem domain. It is also common to set these values based on statistical inference. Within the experiments, statistical inference will be used to determine the correct values.

The same logic can be applied to define membership functions for low and average trust. Furthermore, the same functional form can be applied to the measuring of satisfaction of values and corresponding reputation. For the sake of simplicity, fuzzy logic will only be applied to the final value of trust. Through this application of fuzzy logic, partial membership increases the accuracy in deciding the trustworthiness and subsequent response by the agent.

3.6 Tying It All Together

There are three main components to the new model of trust: reputation, similarity, and self-esteem. The trust mechanism of agent A evaluating the trustworthiness of agent B can be summed up in the following algorithm:

```
Trustworthiness(Agent a, Agent b, Context c)
{
  Self-esteem e1 = a.getPersonalReputation(a, c);
  if (e1 < THRESHOLD_VALUE)
  {
    a.updateValueHierarchy();
  }
  Reputation r1 = a.getPersonalReputation(b, c);
```

```

Reputation r2 = a.getSocialReputation(b, c);
Similarity s1 = a.getValueSimilarity(b);
return fuzzification(ALPHA_1*r1+ALPHA_2*r2+ALPHA_3*s1);
}

```

Within this algorithm, it is seen that the function prototype requires the context of the interaction. Here, context is defined as the different role patterns. If *A* wishes to evaluate *B*'s reputation within the context of an agent assistant, figure 2 dictates that different values must be taken into account when evaluating the reputation of *B* from *A*'s perspective. *A* may wish to change the context of evaluation depending on the composition of *A*'s value hierarchy. For instance, if *A* regards cost-efficiency and capability as very important relative to the other values, it is wise for *A* to consider *B* within the context of a service provider or citizen. It is unwise to consider *B* within the context of an assistant as these values are not manifested within that context. Essentially, *A* would have to evaluate *B*'s reputation based on values that are not as important to *A*. Of course, *A* can always consider *B* within all context by simply considering *B* as fulfilling the agent role. The choice is up to the agent. Within this work, the context is the agent role for the sake of simplicity.

Once the context has been determined, the self-esteem of the agent making the judgement must be assessed. Assessment is made by considering the personal reputation within the selected context. Self-esteem must be first assessed because it can have a great impact on the hierarchy of values held by *A*. Subsequently, the change in value hierarchy alters *A*'s perception of itself along with the perception of others. If the fuzzified self-esteem is considered low, the agent must make changes to its value hierarchy to improve self-esteem. These changes must be authorized by the user or it is conjectured that the user will invalidate the model of trust. Due to the scope of the work, the methodology behind changing the value hierarchy is left for future work.

Once the self-esteem has been evaluated and any necessary changes have been made to the value hierarchy, the agent attains the personal and social reputation of the other agent along with the degree of similarity held between the two agent's value hierarchies. In the event that *A* knows nothing of *B*'s reputation, it can always fall back on the similarity of values held between them. Both reputations are evaluated within the selected context. Personal reputation is based on direct observations while social reputation is based on speculations provided by other agents. A socialization process allows *A* to build a social reputation of *B* by discussing *B*'s actions with others.

The importance of each of these components to determining trustworthiness is expressed through a weighting scheme. In the above pseudocode, *ALPHA_1*, *ALPHA_2*, and *ALPHA_3* must sum to one. The actual weighting scheme behind the balance can be implemented in a number of different ways. For the purposes of this work, the choice of weighting scheme is outside of the scope and has been left for future research. Once a crisp value has been determined for the trustworthiness, it is fuzzified using the technique outlined in Section 3.5.

4 Effectiveness Validation Model

Within human societies, inflation is a common enemy of any economic system. Inflation is defined as a general increase in the cost of consumer goods or a corresponding decrease in the purchasing power of currency [28]. It generally results from an increase in the money supply of an economy. As the money supply of an economy rises, the demand for consumer goods generally rises as well. This increase in demand forces producers to raise prices due to the limited supplies of a good. As a result, the same monetary amount can purchase less than it could before the inflation occurred.

The simulation's objective criterion is the inflation rate. Within an e-commerce framework, market manager agents concern themselves with monitoring the inflation rate amongst agents that are buying and selling within a given market.




Fig. 5. Simulation Participants

Figure 5 denotes the participants within the proposed game. The buying agent B_1 is interested in buying a given quantity of a widget from one of the set of available seller agents $\{S_1, S_2, S_3, \dots, S_n\}$. In this case, B_1 is currently communicating with S_1 as denoted through the use of arrows. S_1 can communicate with all of the sellers as it must know the true price of their widget. A market manager agent MM is denoted by a pentagon and aids in protecting the consumer from deception. The mystery shopper agent MS is denoted as a cross symbol and aids the market manager agent.

Inflation can be established through the creation of artificial demand for a widget within a given business. Artificial demand can be established through the public perception that a particular seller agent's prices and cost of acquisition

are lowest relative to its competitors. The perception results in a false sense of trust between the buyer and seller agent. The consumer agent must believe that the seller agent is indeed trustworthy in order for them to conduct business.

If an agent exercises trust correctly, the harm of deceptive business practices on the society is minimized. Complications arise from the fact that the cost of the widget includes other things besides the cost of the unit itself. It includes fuzzy factors like the cost of ownership, utilization, and maintenance. Furthermore, agents may ask for the price of the competitors' widget. A seller agent can lie about its competitors' prices in any of the different cost factors. The choice to verify the reported price of competitors comes at a cost to the consumer as the consumer agent must communicate with the other competitors in real time.

If an agent has a high suspicion of untrustworthy behavior, it reports its findings to the market manager. The market manager investigates and takes appropriate action. It must independently verify the deception and punish the offender through a reduction in social reputation. Verification occurs through the mystery shopper, denoted as a plus in the presented figure.

4.1 Value Manifestations

The designs of the various agents reflect the manifestation of the values outlined in Section 3. The simulation connects back to the model through an examination of the values and subsequent manifestations within the simulation.

Responsibility is manifested through a buying agent's ability to respond to a user request and make an actual purchase given user expectations. In the event that all selling agents are deemed untrustworthy or unable to satisfy the order, the agent fails to achieve its goal and is perceived as irresponsible by the end-user.

Honesty is manifested through the morality of the selling agent. The defamation detection that has been externalized acts as an indicator of honesty. Service reports are randomly generated and also contribute to the honesty measure. Both components are external to the simulator.

Independence manifests itself through the buying agents willingness to interact with the selling agents without the assistance of other buying agents. Selling agents manifest this through willingness to interact with buyers directly as well.

Obedience is manifested within the buying agents and selling agents through examination of order of goal fulfillment. Within the simulator, the queue of service performance is hard-coded to accept owner processes above agent processes. Hence, obedience becomes a fixed value.

Benevolence is manifested through one buying agent's willingness to perform searches for another buying agent. It is implemented as already discussed.

Capability is manifested within the buying agent as the difference between the widget cost placed upon the agent by the owner and the actual purchasing price of the widget. Within the selling agent, capability is manifested through the selling agent's ability to sell all of the quantity of widgets at the price provided by the producer.

Lastly, cost-efficiency is manifested within the buying agent as the difference in cost to purchase the widgets and the actual savings below the user's expectations. The cost to purchase the widgets is represented as the Euclidean distance between the two agents. The savings are represented as the capability.

5 Analysis

In the previous section the design and implementation of the proposed simulation was presented. This section is devoted to the analysis of the trust model within the proposed simulation. Section 5.1 illustrates and analyzes the impact of deception and the subsequent environment of distrust on the inflation rate. Section 5.2 examines the consistency of public and private perception. Specifically, the issues of identity formation and management are investigated. Section 5.3 considers the stability of the model with respect to trustworthiness ranking. Given the same data for simulations, the analysis seeks to understand the variance on a daily and yearly basis.

The simulation and analysis are based on the trust model introduced in this work using the values and formulas discussed. The values held by the agents are those outlined in Section 3 and manifested as outlined in Section 4.1. The formalizations are strictly adhered to and trust is implemented within the simulation as outlined in Section 3.6. The range of the fulfillment of values is between $[0,1]$ or $[-1,1]$ where appropriate.

There are many variables within the simulations. For simplification, some of these variables have been fixed. The agent population fluctuates between 250 and 500 agents in the appropriate studies. The homogeneity of the value systems held by the agents is a variable needed for the perception congruency analysis conducted in Section 5.2. Within that section, homogeneity is fixed while it is unfixed in all other sections.

The simulator has been implemented using Java 2. A minimal MAS has been implemented that provides the necessary communication support for buyers and sellers. The simulations were carried out on the Linux 2.2.4 and Java 1.1.3 platform. The platform is run on a PC with a Pentium 3 1.6Ghz processor and 256MB of memory.

5.1 Problem Analysis

A healthy environment is defined as one in which the business agents do not engage in deceptive business practices. Buying agents and selling agents can engage each other with full trust. Buying agents do not have to risk a decrease in purchasing power.

Figure 6 denotes the average rate of inflation within a healthy environment. In this case, the net aggregate demand is 4.1% less than the net aggregate supply. Hence, there is a general trend in the annual inflation towards -4.1%. This behavior is a direct consequence of the base-price adjustment formula. Business agents are constrained by the limits of their supply. A primary goal of

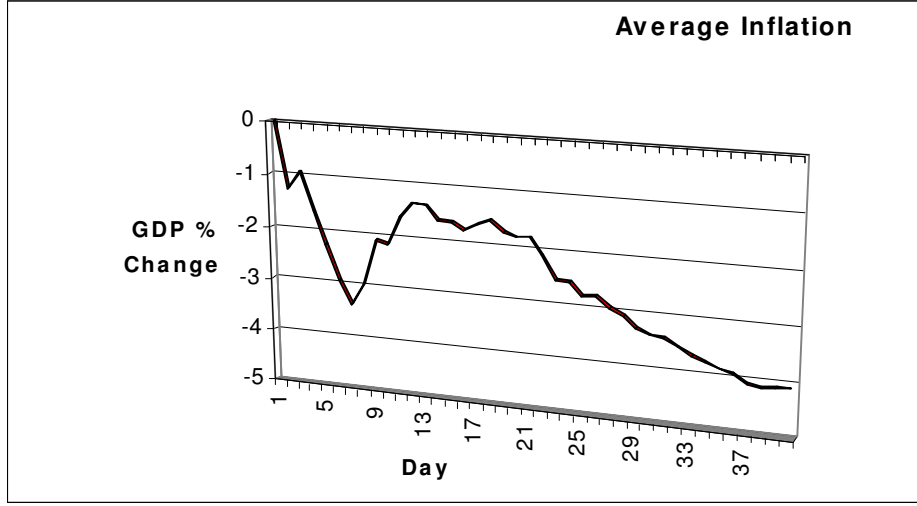


Fig. 6. A healthy system under the finite money supply assumption

the business agent is to eliminate the available supply at the best price possible. If any supply remain at the end of a day, the adjustment formula reduces the current cost of the widget as a function of that supply. Hence, the average price for the widget must fall in the event of decreased demand and subsequent increase in excess supply.

An unhealthy system is one that promotes uncertainty and distrust within. In this context, an unhealthy e-commerce system is one in which selling agents attempt to artificially raise prices in order to maximize profits independent of the current supply. Uncertainty results from the distortion of their own widget price along with that of their competitors. In such an environment, buyers must protect themselves against such gouging through the elimination of uncertainty. In the event that such protection comes at too high a cost, a surprising second factor has been discovered that helps to offer protection.

Figure 7 denotes the inflation within an unhealthy system without any enforced trust mechanisms by the agent. In this scenario, the aggregate demand was approximately 3.0% less than the aggregate supply. 50% of the total selling population is dishonest. Of those that are dishonest, they chose to lie about their competitors' widget prices 75% of the time. When they lie about their competitors, they are also given the freedom to lie about the price of their own widget so long as it makes sense. They can lie about their own widget's price provided they stay in the same ranking of prices as when lying about the competitors. These percentages were arbitrarily set to allow for ease of subsequent analysis.

Based on the previous analysis, one expects the inflation rate to be in proportion to the demand rate under healthy conditions. Surprisingly, the same condition holds true for the case of an unhealthy system under certain condi-

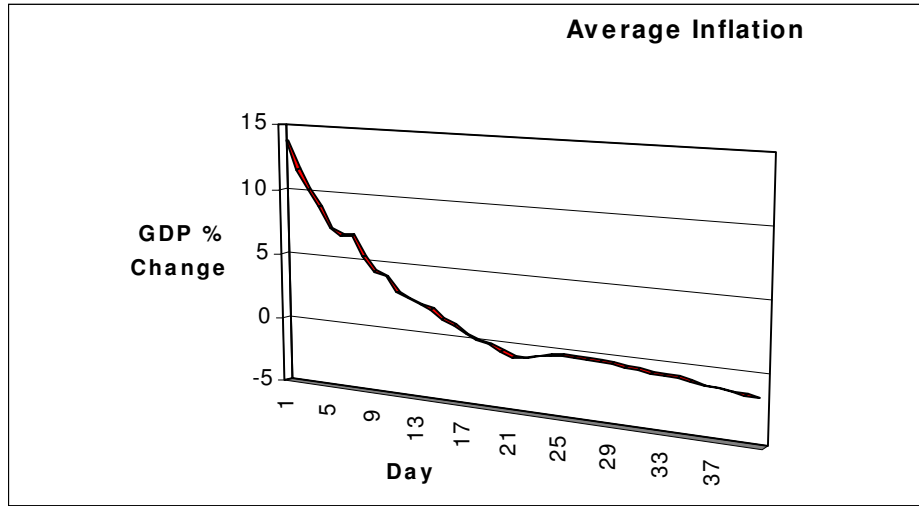


Fig. 7. An unhealthy system with assumptions of finite money supply and blind trust

tions. Initially, there is a surge in the inflation rate as competitors chose to lie about their own prices. The surge in prices translates directly into decreased purchasing power of the agent. In response, agents are able to purchase fewer widgets. Subsequently, the sellers' supplies increase and the supplier is forced to decrease their selling price. Even if they chose to continue lying about the cost of their widget, they will lie on top of a steadily decreasing widget cost.

Ultimately, the economy corrects itself through the seller dependency on supply. This demonstrates that it is best to lie only if the supply chain will not be subsequently altered. Otherwise, the short-term gains of lying will be offset by the long-term loss of demand. Subsequent long-term losses such as loss in reputation (included in the model) further punish chronic price gouging.

Consider the healthy system depicted in Figure 8. As before, this system contains honest sellers that report correct prices every time. However, the aggregate demand is 0.3% greater than the aggregate supply. Essentially, one expects the inflation to equal 0.3% by the end. This time, each agent is given an infinite money supply. Every request by the users will be satisfied at whatever cost necessary.

On any given day, the demand for widgets by any particular agent is random within a certain range. This range is selected so as to enforce a desirable net aggregate demand. As a result, any particular seller should expect to have random fluctuations in demands on any given day. Normally, the fluctuations are met with fluctuations in the widget price.

In the event of an infinite money supply per agent, each fluctuation in price is not met with a corresponding decrease in demand. This makes sense because agent demands are no longer restricted by price. Instead, each surge in demand




Fig. 8. A healthy system under assumption of infinite money supply

has an accumulating effect on price and inflation results. This phenomena is a natural consequence of increasing the money supply to an economy. As a result, the inflation rate is 35% by the end of the simulation.

Figure 9 denotes an unhealthy system in which aggregate demand is 2.4% less than aggregate supply with the precondition of an infinite money supply. In such a system, lying is fully encouraged as there are no safeguards against it and it does not hurt any of the buyers. The inflation rate sores to over 140% within the same time frame as before. The rate of acceleration of inflation has increased dramatically compared to the same system without lying.

It was originally conjectured that inflation itself could be used as a means of detecting lies within these systems under the assumption of a finite money supply. While it has been demonstrated that inflation does initially result from chronic deception, the best means of detection appears to be through the acceleration of inflation under the assumption of an infinite money supply. Under such a system, one expects a natural rate of inflation to occur as suggested in Figure 8. Any inflation acceleration measurements above average may suggest aggregate deception.

Figure 10 demonstrates the effectiveness of the model in preventing this acceleration of inflation under the conditions of an infinite money supply combined with dishonest behavior. As before, the dishonest population is 50% with a propensity to lie 75% of the time about their own widget prices and that of their competitors' widget prices.

Within this scenario, the inflation rate is approximately 45% less than the worst-case scenario of blind trust. Ideally, the inflation rate equals the natural rate of inflation of 35%.




Fig. 9. An unhealthy system with blind trust and infinite money supply assumptions

Explanations for the difference in ideal and actual rate lie within the perception of the agents. As business agents engage in deception, agents detect the lies through the externalized defamation detection mechanism. Subsequently, the perception of honesty of the selling agent becomes more accurate within the deceived agent. The socialization of these deceptive events takes time as agents rely on a peer-to-peer broadcast mechanism. Hence, selling agents can continue to gouge the consumer within a certain time frame before the selling agent's social reputation is seriously impacted. Artificial inflation must result. It is speculated that this artificial inflation rate is also impacted by the propensity to lie, the degree of socialization, and the degree to which honesty is important to agents. The next analysis considers perception in more detail as it relates to public versus private perception.

5.2 Perception Analysis

Ideally, the buying agents of the MAS have an accurate account of the activities of a given selling agent. This account can be summed up through reputation. If buying agents do not have an accurate account of actions, the resulting perception in reputation can lead to unfair gains or losses for the selling agent. A complicating factor is that different agents can have entirely different perceptions of reputation given the same observations of action.

With that in mind, this analysis examines the correctness of public perception. The challenge of verifying the correctness of a perception is two-pronged. Given a reputation value, one must separate the correctness of the observations from the differences in agent opinion about value hierarchy. If both factors are separately examined, one can deduce the validity of agent perception. If one is

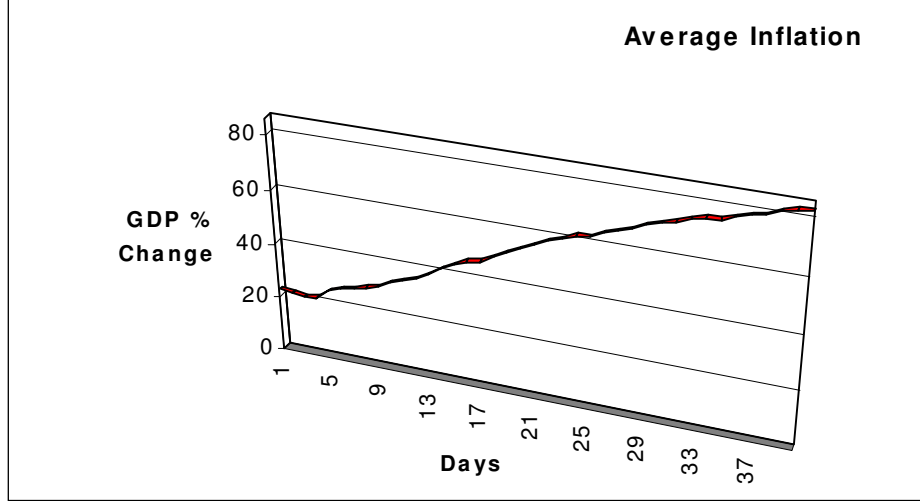


Fig. 10. An unhealthy system with integrated trust model and infinite money supply

just given a final reputation value, it is not possible to separately analyze these two factors.

This is indeed a challenge in the real-world. Given just a person's perception about a topic, one cannot claim the perception is valid or invalid to the observer without knowing both the person's value-system and the actual facts relating to the topic. If the person is mistaken in the actual facts, their perception is misguided. If the person has a different value-system than the observer, the perception is not valid within the observer's context. As a result, it is possible that one person's perception may be incorrect contextually and/or factually.

The only way to work around this problem is to eliminate one of the variables. In this case, the elimination of the context is straight forward. By forcing every agent to adopt a fixed hierarchy of values, the challenge of verifying perception is reduced to verifying the facts. Essentially, the MAS becomes socially homogenous with respect to values. Agents that produce differing perceptions of reputation must result from different perceptions of action alone.

Formally, differences in perception ΔP held by buyer B about seller S are measured through the following equation:

$$\Delta P = \frac{R_S^B}{R_S^S} \quad (14)$$

where R_S^B denotes the reputation ascribed to the seller by the buyer while R_S^S denotes the private reputation or self-esteem of the seller. Ideally, ΔP would equal 1 and buyers will handle sellers appropriately. ΔP of greater than 1 denotes the situation of over-estimation of the seller while ΔP less than 1 denotes an under-estimation of the seller.

Experimentation is carried out within a MAS comprised of 200 buying agents and 25 selling agents. Each agent reports its perception of reputation of every selling agent at the end of every day. All reports for a given selling agent on a given day are averaged. ΔP is translated to an aggregate difference in perception by allowing R_S^B to equal this average. ΔP is observed every day for every selling agent and plotted. Different variables within the model have a direct impact on the aggregate behavior of ΔP . The impact of two variables, honesty and sociability, are analyzed below.

5.2.1 Honesty's Impact on Perception

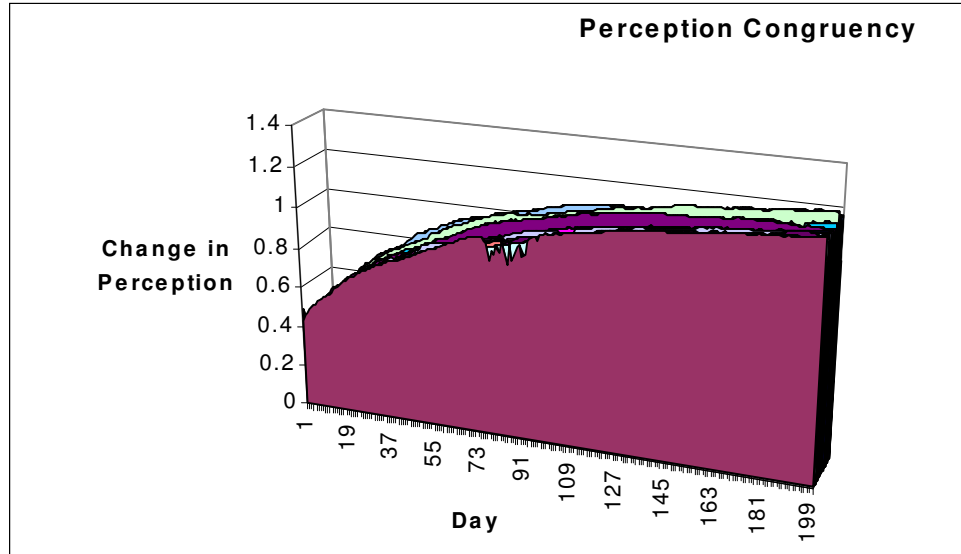


Fig. 11. Honest and Sociable Behavior

Figure 11 shows a MAS with fully honest selling agents and fully sociable buying agents. The degree of sociability of an agent describes its willingness to communicate with its neighbors about a previous transaction. In the case of full sociability, an agent will communicate with all of its neighbors.

As seen, the average initial public perception of a given selling agent can be expected to be less than the perception held by the selling agent. This makes sense as the selling agent is initially an unknown entity and the buying agents make no assumptions about observation of actions. As time progresses, agents socialize with each other and public perception converges as the interactions lead to consistent results. This consistency results from honesty on the part of the seller.

Eventually, public perception approaches the private perception of the selling agent. This results from an increased number of interactions and subsequent observations of action of the selling agent. In this case, the simulation averages a ratio of 1.05 towards the end of the simulation for any given selling agent. A slightly higher estimation of public esteem of the individual may result from short-term decreases in private perception held by the seller. These short-term decreases are natural fluctuations that result from changes in sellers' and buyers' expectation.

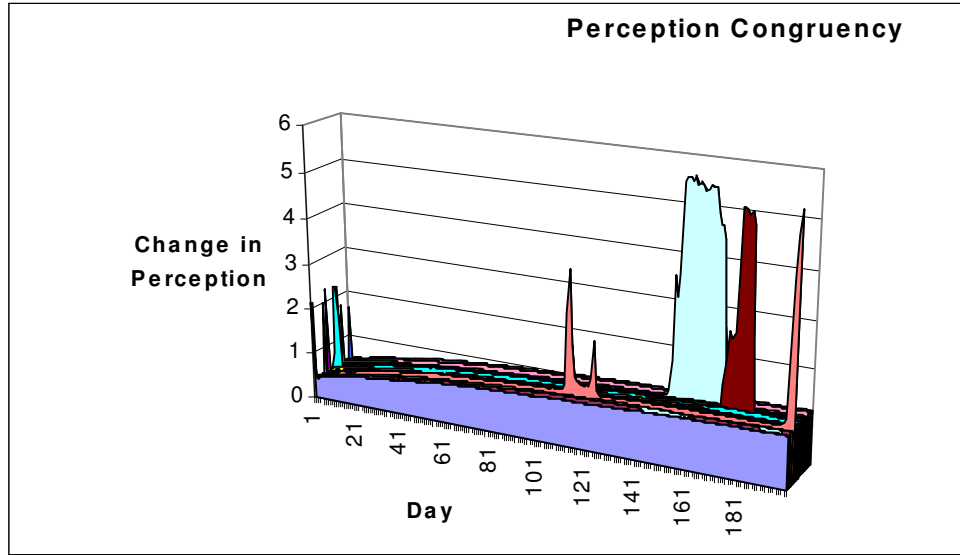


Fig. 12. Dishonest and Sociable Behavior

Figure 12 depicts a similar scenario with a change in honesty of the sellers. In this scenario, spikes in public perception result in severe over-estimation of the reputation of the perceived seller. The seller sees a subsequent gain in flow of business activity while other sellers see a consequential loss. In such a scenario, a highly sociable set of agents quickly correct such spikes through aggregate averaging of speculation. This naturally leads to the question of the impact of sociability on perception.

5.3 Stability Analysis

In order to analyze the stability of the model, one must consider the various potential views of stability. Within this work, stability is connected to the idea of ranking. Each selling agent maintains a certain reputation within the MAS. These agents can be ranked in ascending order of social reputation. The social

reputation can be acquired by averaging the reputation of each seller by each buyer. This is done as before in the perception analysis.

This ranking of sellers can shift on a daily basis as well as on an iterative basis. An iteration is defined as one full run of the simulation. Iterative stability refers to the degree of change of rankings of sellers on a per simulation basis. A high iterative stability implies that agents will not change very much in their overall ranking from simulation to simulation. Meanwhile, daily stability refers to the degree of change of rankings of sellers within any given iteration. A high daily stability implies that agents will not change much in their shift in rankings within any one iteration. It is possible to have any combination of the two types of stability. As an example, a high daily stability with a low iterative stability implies that agents do not have an assured spot in the ranking. However, once they have been given a spot, they remain close to that spot within the iteration.

Due to the random nature of the simulation, descriptive statistics must be used to measure the stability in order to eliminate as much randomness as possible in the data. Stability is measured through an examination of the sample variance of the data. Each data element $\Delta_s(i)$ is a random variable that represents the average change in ranking of seller s within iteration i over all the days within that iteration. If there are n iterations with m sellers, there must be nm data elements present in the set. Figure 13 depicts an example set.

Iteration	Agent				
	1	2	3	...	m
1	1.23	1.43	0.23		6.75
2	1.92	1.2	4.92		0.33
3	1.01	3.21	3.01		4.37
...					
n	1.32	2.12	0.71		4.55

Fig. 13. Example Data of Typical Stability Analysis

5.3.1 Iterative Stability Analysis

Iterative stability analysis focuses on the stability of iterative ranking of agents on a per-iteration basis. If agent A receives an overall ranking of $rank_a$ in itera-

tion i , how likely is it that the same agent will be ascribed a similar ranking in subsequent iterations? The answer can be found through an examination of the sample variance of $rank_a$ over various iterations.




Fig. 14. Average Shifts in Iterative Agent Ranking

Figure 14 denotes the average sample variance experienced by any given agent within a given year or iteration. Essentially, this represents the average sample variance of any row depicted in Figure 13. The range of sample variance is theoretically greater than or equal to 0 and less than 144. On average, the iterative sample variance is 3.88. This means that any given agent can expect a change in their overall ranking within an iteration of plus or minus 3.88 places from where they started on average. In each iteration, the seller population is maintained at a constant of 25. Relatively, this means that each agent will travel in ranking within a spectrum of approximately 31% of the total available holdings. As an example, an agent that places 8th on average within one iteration will rank as high as 12th or as low as 4th in different iterations. Based on the graph, the behavior appears to be consistent and stable.

In a separate set of experiments, the variance experienced per iteration for a specific agent is examined instead of for all agents. The goal is to examine the consistency of the variance rather than the actual values of the variance.

Figure 15 displays the results of the simulation. Rather than examining an average of all agents within a given iteration as before, each agent's average sample variance is measured across all iterations. In this case, Figure 15 demonstrated that the individual average sample variance is less than the collective variance that is considered in the previous analysis. On average the individual




Fig. 15. Average Iterative Shifts Per Agent

variance is 1.30 while the collective variance is 3.88. The range of the iterative shift per agent is approximately 2 while the range of the overall iterative shift is approximately 5. This suggests that the individual variance has a synergistic effect on the variance of the collective. The individual iterative variance appears to be stable and consistent within the fixed conditions of population.

The validity of the fixed conditions raises another question. Does the individual and collective variance remain stable over changes in population? Is this model scalable over a change in environment? The scalability analysis addresses this question.

5.3.2 Daily Stability Analysis

Figure 16 denotes the results of yet another experiment conducted under the same conditions as before. This time, a statistical analysis of an arbitrarily chosen iteration is done. The sample variance within any one iteration for any one agent averages at approximately 6.72. The range of the sample variance is approximately 10. The measured sample variance is higher than the individual variance per iteration experienced by an agent. In other words, the agent experiences high iterative stability and low daily stability.

5.4 Scalability Analysis

Lastly, the intersection of scalability and stability is investigated. The artificiality of the fixed population condition surrounding the stability analysis limits the




Fig. 16. Average Daily Shifts Per Agent

usefulness of the analysis within a more realistic scope. Allowance of population to be variable allows for a test of both scalability and stability.

Figure 17 denotes the results of an iterative stability analysis conducted over changes in the population. As always, buyers and sellers are maintained at a ratio of 10:1. The population pertains only to the buyers. Hence, the real population is 10% higher than the graph reports. The population spans from 250 to 500 in increments of 25. The line of best fit S that describes the variance of the agent population P is defined as:

$$S = 0.0322P - 5.4864 \quad (15)$$

Based on this figure, the iterative variance appears to be linear in behavior. This is ideal behavior as it allows for a consistent and reliable stability of the MAS. Furthermore, the linearity aids as a predictor of the variance given a population. However, the negative Y-intercept of equation 15 indicates that the slope is a very sensitive approximation. As P approaches 0, S approaches -5.4864 . Small sample changes in the slope will result in larger changes in the constant.

Figure 17 denotes the results of separate simulations carried out under the conditions of daily variance. Once again, linear behavior is exhibited within the changes of social ranking on a daily basis:




Fig. 17. Scalability of Variance

$$S = 0.0172P - 3.1249 \quad (16)$$

Equation 16 is the best-fit of the sample points provided. Once again, a negative Y-intercept suggests that the measurement of slope is an approximation that is sensitive to the samples taken. A small change in the measurement of the slope will result in a larger change in the constant.

6 Conclusions

The model of trust presented strives to account primarily for two differing properties of trust. Trust is objective in the sense that it can be based solely on fact. It can be highly rational and clear in specific cases. Trust can also have emotional components that are wrought with vagueness and uncertainty. This model combines both the observation and uncertainty components through reputation, values, familiarity, and self-esteem.

Reputation becomes both objective and subjective through the use of interactions and speculations. Interactions account for the objective component while speculations account for the subjective or unknown component.

Values account for the subjective nature of the rationalization component of trust. It sounds odd to say that rationalization is indeed subjective. It is true according to value theory. Essentially, values act as a road map in making decisions. The importance of different values are learned and entirely subjective. With that in mind, the decision of trust cannot be broken down to one rationalization but many potential rationalizations. Each person has their own separate rationalization that must be accounted for. Otherwise, the rationalization is contextually wrong and invalidated by the person. Computer systems that exhibit trust on behalf of users must take this into account or risk invalidation.

Familiarity is exhibited through the Euclidean distance between two vectors that describe the value hierarchies of two users. Within the implementation, it is reduced to a speculative component of reputation. Familiarity is a necessary precondition to trust.

Lastly, self-esteem is represented as the reputation an agent has with itself. Self-esteem acts as a barometer of the value hierarchy. In the event that low self-esteem results, corrective measures can be taken. The value hierarchy can be adjusted. This is forcing the agent to adjust its own perception of the world and of itself. As a second measure, the actions taken by the agent can be altered so as to fulfill the values reflected in the value hierarchy.

Within the previous work conducted on reputation, it was claimed that reputation could not be universalized due to the dependency on roles. Through the integration of value theory into the trust model, a set of values have been established that are independent of the MAS containing them. These values now manifests themselves through roles that are universal. Hence, the challenge of universalizing reputation has been overcome through the integration of value theory.

Lastly, one of the great challenges to computer science lies in the integration of computers and the culture they reside in. Currently, the internationalization of software focuses mostly on esthetic aspects of programs. These include color, fonts, language, window size, and orientation. Agents fail to take into account culturally-specific values when rationalizing a decision. Most rationalization processes are fixed and force one particular culture or set of beliefs. Internationalization of agents must take external values into account or risk invalidation by the user. This model focuses on validation by the end user.

6.1 Quantitative Aspects

Originally, the intention of the game was to allow an objective measurement of the effectiveness of the trust model through the inflation rate. Surprisingly, the analysis reveals that underlying macroeconomic principles take care of deceptive business practices without aid. Once the assumption of a finite money supply is lifted, the second derivative or acceleration of inflation does indeed need assistance by the model to curb the effects of deception. Without aid, the inflation approaches 140%. With the model, inflation reaches approximately 80% or 50% above the naturally occurring rate of inflation. While this is less than ideal, it demonstrates the effectiveness of the model in addressing the acceleration of inflation in an e-commerce environment.

Next, the analysis of the perceptual symmetry between a selling agent and the buying agents at large was conducted. The model demonstrates many phenomena that mimic real life in terms of identity formation. It is shown that as time progresses, the general public perception of reputation becomes more accurate to that of the private perception. Initially, the buying agents have a lower esteem of the individual than does the individual. This makes sense as the individual is an unknown element to the public that must earn a reputation through

repeated interactions. Eventually, the public and private perception converge provided that deception does not occur.

In the event of deception, the selling agent is attempting to raise the public's esteem higher than the seller's own self-esteem. The model demonstrates that this does indeed occur. Socialization impacts the process through a magnification of the effects of deception. The degree of socialization directly influences the rate at which the public's esteem of the seller is influenced.

The impact of honesty on public perception was also examined to verify the sensibility of the model. The honesty of the sellers has a similar influence on the rate at which the public's esteem of the seller is influenced. This makes sense because agents that are successfully dishonest will be able to directly influence the rate of increasing public perception. The model reinforces this common-sense notion.

The practicality of the model is considered through an examination of the stability and scalability. The stability analysis reveals that the model exhibits high iterative stability and lower daily stability. Indirectly, it is concluded that the trust described by the model exhibits the same variance. Scalability addresses the issue of stability over changes in population. The scalability analysis reveals that the model is scalable over changes in agent population. The model exhibits linear behavior in iterative and daily variance over changes in population from 250 to 500 agents.

References

1. A. Abdul-Rahman and S. Hailes. Relying on trust to find reliable information. In *1999 International Symposium on Database, Web and Cooperative Systems (DWA-COS'99)*. University College London, 1999.
2. A. Abdul-Rahman and S. Hailes. Supporting trust in virtual communities. In *Hawaii International Conference on System Sciences*. University College London, 2000.
3. J. Ball-Rokeach and W. Loges. *Value Theory and Research*. Macmillan, 2000.
4. R. Bhattacharya, T. Devinney, and M. Pillutla. A formal model of trust based on outcomes. *Academy of Management Review*, 23(3):459–472, 1998.
5. E. Bitting. Defamation within multiagent systems. Master's thesis, University of New Brunswick, December 2002.
6. N. Branden. *The Six Pillars of Self Esteem*. Bantam Books, 1995.
7. J. Brayton. The social shaping of virtual reality: A feminists analysis. Master's thesis, Queen's University, Kingston, Ontario, Canada, May 1995.
8. R. J. Brym. *New Society, Sociology for the 21st Century*. Harcourt Brace and Company, Canada, 1998.
9. J. Carter, E. Bitting, and A. Ghorbani. Reputation formalization within information sharing multiagent architectures. *Computational Intelligence*, 2(5):45–64, 2002.
10. M. S. Corey and G. Corey. *Groups - Process and Practice Sixth Edition*. Brooks/Cole Thomson Learning, Canada, 2002.
11. S. Dollinger and F. Leong. On traits and values: With special reference to openness to experience. *Journal of Research in Psychology*, 30, 1996.

12. L. Duhaime. Canadian defamation law.
13. G. C. for Free Expression. Defining defamation: Principles on freedom of expression and protection of reputation.
14. E. Goffman. *The Presentation of Self in Everyday Life*. Horizon Scientific Press, 1959.
15. M. Gordon. *The Concise Oxford Dictionary of Sociology*. Oxford University Press, 1994.
16. P. Hwang. Properties of trust: An analytical view. *Organizational Behavior and Human Decision Processes*, 69(1):67–73, January 1997.
17. J. S. Hyde, J. D. DeLamater, and E. S. Byers. *Understanding Human Sexuality Canadian Edition*. McGraw-Hill Ryerson, Toronto, Ontario, 2001.
18. F. Kluckhohn. *Values and Value Orientation in the Theory of Action*. Harper and Row, 1950.
19. J. Lewis and A. Weigert. Trust as social reality. *Social Forces*, 63(4):976, 1985.
20. N. Luhmann. *Trust and Power*. John Wiley, 1979.
21. R. Mayer, J. Dave, and F. Schoorman. An integrative model of organizational trust. *Academy of Management Review*, 20:709–734, 1995.
22. A. M. Mohammed. *Benevolent Agents*. PhD thesis, University of South Carolina, 2000.
23. U. of International Associations.
24. P. Pavlou and S. Ba. Does online reputation matter? an empiricial investigation of reputation and trust in online auction markets. In *Proceedings of the Americas Conferece on Information Systems*. Association for Information Systems, Long Beach, 2000.
25. N. Rescher. *Introduction to Value Theory*. Prentice-Hall Inc., 1969.
26. M. Rokeach. *The Nature of Human Values*. The Free Press, 1973.
27. D. Rousseau, B. Sitkin, R. Burt, and C. Cameber. Not so different after all: A cross-discipline view of trust. *Academy of Management Review*, 23(3):393–404, 1998.
28. J. Sloman and K. Norris. *Macroeconomics*. Addison-Wesley, 2nd edition, 1999.
29. R. L. Solso. *Cognitive Psychology Sixth Edition*. Allyn and Bacon, Needham Heights, MA, 2001.
30. N. Szabo, 1996.
31. W. Weiten. *Psychology Themes and Variations, Briefer Version*. Brooks/Cole Publishing Company, Pacific Grove, CA, 1997.
32. J. Yen and R. Langari. *Fuzzy Logic Intelligence, Control, and Information*. Prentice Hall, Upper Saddle River, New Jersey, 1999.