An Analytic Approach to Reputation Ranking of Participants in Online Transactions

Aleksandar Ignjatovic, Norman Foo and Chung Tong Lee
School of Computer Science and Engineering
University of New South Wales and
National ICT Australia (NICTA)
Sydney, Australia
e-mail: {ignjat, norman, ctlee}@cse.unsw.edu.au

Abstract

In our setup, agents from a community interact in pairwise transactions across discrete time. Each agent reports its evaluation of another agent with which it has just had a transaction to a central system. This system uses these time-sequences of experience evaluations to infer how much the agents trust each another. Our paper proposes rationality assumptions (also called postulates or constraints) that such inferences must obey, and proceeds to derive theorems implied by these assumptions. A basic representation theorem is proved. The system also uses these pairwise cross-agent trustworthiness to compute a reputation rank for each agent. Moreover, it provides with each reputation rank an estimate of the reliability, which we call weight of evidence. This paper is different from much of the current work in that it examines how a central system which computes trustworthiness, reputation and weight of evidence is constrained by such rationality postulates.

1 Introduction

With the rapid increase of commerce and other online transactions, it has become advantageous to have methods for evaluation of the behavior of agents involved. The notions that address how agents assess other agents, and how overall reputations arise, derive from cognitive social science. The primary reference for this background is the foundational work [1] to which we refer the reader; however, in our paper we will explain those notions pertinent to the issues we raise. This paper is different in spirit of much of the ongoing work on reputation of which [5] is a good example. In [5] the main method for evaluating performance of different approaches is by simulation. Our paper however sets out postulates that are intuitively appealing for any such system, and from them we derive representation theorems.

An agent's assessment of another agent with which it had a transaction is its *experience evaluation*; such evaluation may vary for different transactions when they repeatedly interact over time. Thus, an agent A_i will have an *experience sequence* $\vec{s}_{ij}$ of its transactions with agent A_j, usually evaluating agent A_j for the transaction at time instant t with $e_{ij}(t)$ on some scale. Experience sequences are the bases of an agent's judgement of another agent's *trustworthiness*. The trustworthiness can be seen as an aggregation over time of all individual past experiences into a single quantity on some scale of ranking. Such judgements of trustworthiness of an agent by other agents from the same community can be further aggregated into an easily understood numeric or qualitative rank. We call this *community-wide judgement* on the trustworthiness of an agent its *reputation* within the community.

Both aggregation over time and aggregation over the community of agents are needed to provide a public assessment of the reputation of each agent involved in pairwise transactions. In some models in the literature agents combine their experiences over transactions and report their trustworthiness assessments of other agents. We argue that this is **not** a good approach and that agents should report **only** their experiences in each individual transaction. The reason is that the assessments of agents should be as objective as possible. Thus, agents should provide the most uncompounded, raw data, to avoid as much as possible individual bias of agents that might occur if it is left up to the agents to aggregate individual experiences in unspecified way. In our approach, these two levels of aggregation of individual experiences are performed by a central mediating *agency*. This central agency can maintain the privacy of the experience evaluations of individual transactions given by the agents while publishing only its final results of aggregations, which further increases the robustness of the system against possible manipulative reporting. Details of the aggregation methods used by the agency may also be published.

Fair conditions, also called rationality assumptions, that constrain how an experience sequence $\vec{s}_{ij}$ is related to the consequent trustworthiness assessment have been investigated in [4] and [2]. These assessments among agents up to a time instant t are further aggregated into a community evaluation of each agent at t, which we call *reputation* for agent A_j at t.

Some well-known and widely used systems already give some data to help assess the reputation of an agent. For example, the *eBay* makes available for each agent the percentage of positive reports it receives from other agents, called the *"Feedback"*. The report [6] has a comprehensive survey of many alternative approaches to representing, assessing, computing and updating agent trustworthiness and reputation.

In this paper we provide a rigorous analysis of the problem of ranking the trustworthiness and reputation of agents involved in transactions, and in assessing reliability of such a rank. We provide some rationality assumptions that such ranking methods should satisfy. While assumptions that are essential for our analysis are persuasive, much of our technical development can be modified for somewhat different setups that do not meet all of our assumptions.

1.1 Setup

Assume that N agents $\{A_i : 1 \leq i \leq N\}$ are involved in some pairwise transactions over discrete time moments; thus, we identify time instants with natural numbers. If at an instant k an agent A_i has a transaction with another agent A_j, it records an evaluation of its *experience* for that transaction, $e_{ij}(k)$. Over time up to the instant n, agent A_i has produced a sequence of experience evaluations which we identify with the *partial function* $\vec{s}_{ij} = \{(k, e_{ij}(k)) : k \in D(\vec{s}_{ij})\}$, with the domain $D(\vec{s}_{ij})$ as the collection of time instants $k \leq n$ when agent A_i was engaged in a transaction with A_j; we assume that the "granularity" of time intervals is sufficiently fine so that any two agents at any instant can be involved in at most one transaction. Each agent A_i produces one such sequence for every agent A_j that it has had transactions with. If $k \in D(\vec{s}_{ij})$, then we set $\vec{s}_{ij}[k] = e_{ij}(k)$. If no confusion can occur, we drop indices i, j and write $e(k)$ and $\vec{s}$ for $e_{ij}(k)$ and $\vec{s}_{ij}$ respectively.

After a transaction at moment k, the values $e(k)$ are reported to a *central agency* and the agency computes a trustworthiness assessment $T(\vec{s})$ on the basis of the sequence $\vec{s}$ of *all* experiences of agent A_i with agent A_j up to the instant n. We call this assessment *trust rank* for short. The system then uses $T(\vec{s})$ for all pairs of agents A_i, A_j involved in a transaction at some instants $k \leq n$, to assign to each agent A_j a *reputation rank* $\rho_j(n)$, valid for the moment n, that reflects adequately the trust ranks that other agents ascribed to agent A_j, as well as an estimate of the *weight of evidence* $W_j(n)$ for such rank. While it might be the case that only the ranking system can see the experience reports of individual agents involved, the table of ordered pairs $\mathcal{R}_n = \{(\rho_j(n), W_j(n)) : j \leq N\}$ of the assigned reputation ranks and the evidence weights for those ranks for each instant n is made immediately available to all agents.

Rationality assumptions that should constrain trust rank function $T(\vec{s})$ are proposed, some of which resemble those in [4]. We also discuss how these individual trust ranks should then be aggregated in the community of agents to obtain a community reputation of each agent under similar assumptions.

In the literature it is often assumed that evaluation of experience can be positive, neutral or negative, with positive values increasing the trust rank, negative decreasing it and neutral leaving it unchanged. We argue that such model does not capture our intuitive notion of trust. In particular, it allows an agent to maintain initial high trust simply by maintaining performance that renders experience of other agents neutral.

Let us elaborate this point further by assuming that an agent A_i is dealing at an instant $n+1$ with another agent A_j with which it had a sequence of past experiences $\vec{s}$ with the domain $D(\vec{s}) \subseteq \{1, \ldots, n\}$. Thus, prior to the transaction at the moment $n+1$, A_j's trust rank with A_i is $\tau = T(\vec{s})$. A more satisfactory model should have means of producing the *expected* experience level $\varepsilon(\tau)$ that an agent should expect to get from an agent with trust rank τ. Let $\vec{s}\,'$ denote the sequence of experiences obtained by extending the sequence of experiences $\vec{s}$ with experience $e(n+1)$ at instant $n+1$, i.e., $\vec{s}\,' = \vec{s} \cup \{(n+1, e(n+1))\}$. We argue that a trust rank function $T(\vec{s})$ should satisfy

$$T(\vec{s}\,') \lesseqqgtr T(\vec{s}) \Leftrightarrow e(n+1) \lesseqqgtr \varepsilon(T(\vec{s})).$$

In this way, the trust rank $T(\vec{s}\,')$ of agent A_j after the transaction at instant $t+1$ will be larger than $T(\vec{s})$ just in case its performance $e(n+1)$ evaluated by agent A_i at instant $n+1$ is higher than what one would expect from an agent with trust rank $T(\vec{s})$.

The simplest way of achieving this is to measure the reputation rank by the value of the expected experience level. This means that the experience and the reputation rank should be measured on the *same scale*, with $\varepsilon(\tau) = \tau$, and that the following should hold:

$$T(\vec{s}\,') \lesseqqgtr T(\vec{s}) \Leftrightarrow e(n+1) \lesseqqgtr T(\vec{s}).$$

We will assume that an experience evaluation $e(k)$ is a real number in a fixed range $[0, M]$, $M > 0$. If $e(k) = M$, this means the reporting agent A_i regards the quality of the transaction with A_j at moment k is the best, while a report of 0 means the worst.

2 Aggregation in Time

We start by examining methods for aggregation in time. In previous section, we introduced the notion of *the trust rank* computed on the basis of the experience sequence $\vec{s}$, denoted by $T(\vec{s})$. Should this function be recursive i.e., $T(\vec{s}') = U(T(\vec{s}), e(n+1))$ for some *trust update* function $U(x, y)$, evaluation will be very efficient. We will discuss this aspect at the end of this section. Now, let us list axioms that such a trust rank function T, which aggregates over time individual experiences in a sequence of experiences $\vec{s}$ of an agent A_i with an agent A_j, should satisfy. These axioms attempt to delineate the "rationality" of T.

2.1 Axioms for Aggregation in Time

(T1) Shift Invariance. Let $\vec{s}^+$ be the forward unit shift of the sequence $\vec{s}$, i.e., let $\vec{s}^+ = \langle (k+1, e(k)) : k \in D(\vec{s}) \rangle$; then $T(\vec{s}^+) = T(\vec{s})$. □

This property formalizes our intuition that the trust rank of a sequence of experiences $\vec{s}$ depends only on the pattern of past individual experiences. In other words, the way how the trust rank of A_j with A_i is obtained by aggregation of the past experiences of A_i with dealing with A_j does not change with time and is obtained for all pairs of agents in the same way.

(T2) Time Averaging. Let $\vec{s}$ be any sequence of experiences; then $\min(\vec{s}) \leq T(\vec{s}) \leq \max(\vec{s})$. □

Thus, if $\vec{s}$ is a constant sequence, then $\min(\vec{s}) = \max(\vec{s}) = T(\vec{s})$. In particular, if $D(\vec{s})$ contains a single instant k, then $T(\vec{s}) = \vec{s}[k] = e(k)$. Intuitively, the trust rank $T(\vec{s})$ obtained on the basis of the sequence of experiences $\vec{s}$ is certain form of an average of trust included in the sequence $\vec{s}$.

(T3) Consistency. Let $\vec{s}_1$ and $\vec{s}_2$ be any two experience sequences for the same pair of agents over two disjoint time domains $D(\vec{s}_1) \cap D(\vec{s}_2) = \emptyset$, such that $T(\vec{s}_1) = T(\vec{s}_2)$; then $T(\vec{s}_1) = T(\vec{s}_1 \cup \vec{s}_2) = T(\vec{s}_2)$. □

This formalizes our intuition that if the trust rank of two sequences are the same, then joining them together should also result in a sequence with the same trust rank.

(T4) Discounting. Assume that $0 \leq e < E \leq M$ and let $\vec{s}_1 = \{(1, e), (2, E)\}$ and $\vec{s}_2 = \{(1, E), (2, e)\}$; then $T(\vec{s}_1) \geq T(\vec{s}_2)$. □

Discounting formalizes our intuition that older experiences of an agent A_i dealing with A_j cannot be more important than the more recent ones, because if we compare both sequences $\vec{s}_1$ and $\vec{s}_2$ with $\vec{s} = \{(1, e), (2, e)\}$, increasing the more recent experience increases the trust rank at least as much as increasing the older experience.

We feel that every reasonable trust rank function must satisfy the above axioms. Intuitive justification of the next axiom given below is perhaps somewhat less than universal; however, it greatly simplifies technical matters, without imposing any undue restrictions.

(T5) Linearity. Let α, β be real numbers, and $\vec{s}_1$ and $\vec{s}_2$ two experience sequences such that $D(\vec{s}_1) = D(\vec{s}_2) = D$ and $0 \leq \alpha \vec{s}_1[k] + \beta \vec{s}_2[k] \leq M$ for all $k \in D$. Further, the sequence $\vec{s} = \alpha \vec{s}_1 + \beta \vec{s}_2$ is defined on the same domain D so that for all $k \in D$, $\vec{s}[k] = \alpha \vec{s}_1[k] + \beta \vec{s}_2[k]$. Then $T(\vec{s}) = \alpha T(\vec{s}_1) + \beta T(\vec{s}_2)$. □

Linearity implies that if we increase all experiences in an experience sequence by say 10% and if these values are still below the maximal value M, then the trust rank of that sequence will increase by 10% as well. Also, let $\vec{\Delta}$ be an arbitrary sequence or real numbers with the domain $D(\vec{\Delta}) = D(\vec{s})$. Then for every sequence $\vec{s}$ if we change the values $\vec{s}[k]$ into values $\vec{s}[k] + \vec{\Delta}[k]$, and if these values are within the range $[0, M]$, the trust rank of the new sequence is obtained as $T(\vec{s}) + T(\vec{\Delta})$, i.e., the change of the trust rank does not depend on the starting sequence values, but only on the sequence of changes $\vec{\Delta}$. In particular, improving the experience at any particular instant k in the domain of the sequence $\vec{s}$ for a certain amount increases the overall trust rank of the sequence as much as degrading the same experience at the same instant for the same amount decreases the trust rank. Thus, these features, equivalent to linearity, are reasonable properties for a trust rank function. However, it is not hard to construct quadratic or higher-degree trust rank functions that satisfy all other axioms. Some of our arguments can be extended to such functions as well, albeit with considerable technical complications and, we feel, with doubtful gain.

2.2 Canonical Form of Trust Rank Functions

In this subsection we use (T1) through (T4) to obtain a canonical form for the trust rank function T that maps experience sequences into trust ranks. The main theorem constructs this form, and its corollaries exhibit intuitively appealing properties of that canonical T.

We define $\|\vec{s}\| = \max(D(\vec{s})) - \min(D(\vec{s})) + 1$. Thus, if $\nu(\vec{s})$ is the sequence obtained by translating the sequence $\vec{s}$ such that $\min(D(\nu(\vec{s})) = 1$, then $\max(D(\nu(\vec{s})) = \|\vec{s}\|$. The Shift Invariance Axiom implies that $T(\nu(\vec{s})) = T(\vec{s})$; thus, it is enough to consider trust ranks of sequences such that $1 \in D(\vec{s})$.

The next lemma reveals a form for trust rank functions analogous to the response of a linear system.

Lemma 2.1 *Let n be an integer and $I_n = \{i : 1 \leq i \leq n\}$. Further, let $D \subseteq I_n$ be such that $\{1, n\} \subseteq D$. For a sequence $\vec{s}$ with $D(\vec{s}) = D$, if the trust rank function T satisfies T1 through T5, there exist positive numbers w_k^D such that*

$$\textstyle\sum_{k \in D} w_k^D = 1 \tag{1}$$

$$T(\vec{s}) = \textstyle\sum_{k \in D} w_k^D \vec{s}[k]. \tag{2}$$

Proof. For each $k \in D$, we define the experience sequence $\vec{s}^{\,D}_k$ such that $D(\vec{s}^{\,D}_k) = D$; $\vec{s}^{\,D}_k[k] = 1$ and $\vec{s}^{\,D}_k[m] = 0$ for all $m \in D, m \neq k$. Then we have $\vec{s} = \sum_{k\in D} \vec{s}[k]\ \vec{s}^{\,D}_k$. Since T is linear, $T(\vec{s}) = \sum_{k\in D} \vec{s}[k]\ T(\vec{s}^{\,D}_k)$. Consider the sequence $\vec{1}_D$ such that for all $i \in D$, $\vec{1}_D(i) = 1$. $T(\vec{1}_D) = \sum_{k\in D} T(\vec{s}^{\,D}_k)$. By Averaging Axiom, $T(\vec{1}_D) = 1$; thus, $\sum_{k\in D} T(\vec{s}^{\,D}_k) = 1$, and for all $k \in D$ we can take $w^D_k = T(\vec{s}^{\,D}_k)$. □

Corollary 2.2 (Uniform Continuity) *For every δ and every two experience sequences $\vec{s}_1, \vec{s}_2$ with $D(\vec{s}_1) = D(\vec{s}_2) = D$ and $|\vec{s}_1[k] - \vec{s}_2[k]| < \delta$ for all $k \in D$, then also $|T(\vec{s}_1) - T(\vec{s}_2)| < \delta$.*

Proof. We first note that $\vec{s}_1 - \vec{s}_2$ might not be an experience sequence because some of the values $(\vec{s}_1 - \vec{s}_2)[k]$ might be negative. However, function T, by linearity, naturally extends via (2) to such sequences as well, and $T(\vec{s}_1 - \vec{s}_2) = T(\vec{s}_1) - T(\vec{s}_2)$. Thus,

$$|T(\vec{s}_1 - \vec{s}_2)| \leq \sum_{k\in D} (w^D_k |\vec{s}_1[k] - \vec{s}_2[k]|) < \delta \left(\sum_{k\in D} w^D_k\right) = \delta.$$

□

The following corollary can be proved in the same manner.

Corollary 2.3 (Point-wise monotonicity) *Let $\vec{s}_1, \vec{s}_2$ be two experience sequences with $D(\vec{s}_1) = D(\vec{s}_2) = D$. If $\vec{s}_1[k] \geq \vec{s}_2[k]$ for all $k \in D$, then $T(\vec{s}_1) \geq T(\vec{s}_2)$.*

Theorem 2.4 (Representation Theorem) *Let T satisfy the axioms T1 through T5. Then there exists $q \geq 1$ such that for all $\vec{s}$ the following canonical representation of $T(\vec{s})$ holds:*

$$T(\vec{s}) = \sum_{k\in D(\vec{s})} \frac{q^{k-1}}{\sum_{l\in D(\vec{s})} q^{l-1}}\, e(k). \tag{3}$$

Proof. Note that, by Lemma 2.1, for some $w^{I_2}_1, w^{I_2}_2$ such that $w^{I_2}_1 + w^{I_2}_2 = 1$ we have $T(\{(1,1),(2,0)\}) = w^{I_2}_1$ and $T(\{(1,0),(2,1)\}) = w^{I_2}_2$. By the Axiom of Discounting, $w^{I_2}_1 \leq w^{I_2}_2$. Thus, $w^{I_2}_1 \leq 1/2$ and so for some $q \geq 1$, $T(\{(1,1),(2,0)\}) = 1/(1+q)$ and consequently $T(\{(1,0),(2,1)\}) = w^{I_2}_2 = 1 - w^{I_2}_1 = q/(1+q)$.

We first prove the statement for all D such that $D = I_n = \{1,\dots,n\}$ for some $n \geq 3$, and for q as above. Using Lemma 2.1, $T(\vec{s}) = \sum_{k=0}^{n} w^{I_n}_k \vec{s}[k]$. Fix $m < n$ and consider a set of 3 experience sequences $\vec{s}^{\,mm'}$, $\vec{s}^{\,\bullet mm'}$ and $\vec{s}^{\,\star mm'}$ with $D(\vec{s}^{\,mm'}) = I_n, D(\vec{s}^{\,\bullet mm'}) = \{m, m+1\}, D(\vec{s}^{\,\star mm'}) = I_n \setminus D(\vec{s}^{\,\bullet mm'})$. We define $\vec{s}^{\,mm'}$ as follow:

$$\vec{s}^{\,mm'}[k] = \begin{cases} 1 & \text{if } k = m \\ 0 & \text{if } k = m+1 \\ 1/(1+q) & \text{otherwise ;} \end{cases}$$

while $\vec{s}^{\,\bullet mm'} = \{(m,1),(m+1,0)\}$ and $\vec{s}^{\,\star mm'}[k] = 1/(1+q)$ for all $k \in D(\vec{s}^{\,\star mm'})$. Then by the Shift Invariance Axiom, $T(\vec{s}^{\,\bullet mm'}) = T(\{(1,1),(2,0)\}) = 1/(1+q)$. By the Averaging Axiom, $T(\vec{s}^{\,\star mm'}) = 1/(1+q)$. Thus, by the Consistency Axiom, $T(\vec{s}^{\,mm'}) = T(\vec{s}^{\,\bullet mm'} \cup \vec{s}^{\,\star mm'}) = 1/(1+q)$. Since by Lemma 2.1 $\sum_{k\in I_n} w^{I_n}_k = 1$, we get $T(\vec{s}^{\,mm'}) = w^{I_n}_m + (1 - w^{I_n}_m - w^{I_n}_{m+1})/(1+q) = 1/(1+q)$, which implies $w^{I_n}_{m+1} = w^{I_n}_m\, q$. Consequently, for all $k \in I_n$, $w^{I_n}_k = q^{k-1}/(1+\ldots+q^{n-1}) = q^{k-1}(q-1)/(q^n-1)$.

Consider now another set of sequences $\vec{s}^{\,1n}$, $\vec{s}^{\,\bullet 1n}$ and $\vec{s}^{\,\star 1n}$ with $D(\vec{s}^{\,1n}) = I_n, D(\vec{s}^{\,\bullet 1n}) = \{1, n\}, D(\vec{s}^{\,\star 1n}) = I_n \setminus D(\vec{s}^{\,\bullet 1n})$. Let this time $\vec{s}^{\,\bullet 1n} = \{(1,1),(n,0)\}$; then there exists $r > 0$ such that $T(\vec{s}^{\,\bullet 1n}) = 1/(1+r)$. For all $k \in D(\vec{s}^{\,\star 1n})$, $\vec{s}^{\,\star 1n}[k] = 1/(1+r)$. We define $\vec{s}^{\,1n}$ as follow:

$$\vec{s}^{\,1n}[k] = \begin{cases} 1 & \text{if } k = 1 \\ 0 & \text{if } k = n \\ 1/(1+r) & \text{otherwise ;} \end{cases}$$

Then again by the Averaging and the Consistency Axioms, $T(\vec{s}^{\,1n}) = T(\vec{s}^{\,\bullet 1n}) = T(\vec{s}^{\,\star 1n}) = 1/(1+r)$. This implies that

$$T(\vec{s}^{\,1n}) = \frac{1}{\sum_{k=1}^{n} q^{k-1}} + \frac{1}{1+r}\,\frac{\sum_{k=2}^{n-1} q^{k-1}}{\sum_{k=1}^{n} q^{k-1}} = \frac{1}{1+r},$$

which is easily seen to yield $r = q^{n-1}$. This means that for every $\vec{s}$ such that $D(\vec{s}) = \{1, n\}$ we have

$$T(\vec{s}) = \frac{\vec{s}[1]}{1+q^{n-1}} + \frac{\vec{s}[n]\, q^{n-1}}{1+q^{n-1}}.$$

Finally we consider an arbitrary $D \subseteq \{1,\dots,n\}$ and an arbitrary $\vec{s}$ such that $D(\vec{s}) = D$. Then $T(\vec{s}) = \sum_{k\in D} w^D_k \vec{s}[k]$. To complete the proof of the theorem it is enough to show that if $l, m \in D$ and $l < m$, then $w^D_m / w^D_l = q^{m-l}$. Consider a similar set of sequences $\vec{s}^{\,lm}$, $\vec{s}^{\,\bullet lm}$ and $\vec{s}^{\,\star lm}$. Let $\vec{s}^{\,\bullet lm} = \{(l,1),(m,0)\}$. By the Shift Invariance, $T(\vec{s}^{\,\bullet lm}) = T(\{(1,1),(m-l+1,0)\}) = 1/(1+q^{m-l})$. Again, we set $\vec{s}^{\,lm}[l] = 1$, $\vec{s}^{\,lm}[m] = 0$ and $\vec{s}^{\,lm}[k] = 1(1+q^{m-l})$ for all $k \in D \setminus \{l, m\}$. Likewise, $\vec{s}^{\,\star lm}[k] = 1/(1+q^{m-l})$ for all $k \in D \setminus \{l, m\}$. Once more, we have $T(\vec{s}^{\,lm}) = T(\vec{s}^{\,\bullet lm} \cup \vec{s}^{\,\star lm}) = T(\vec{s}^{\,\star lm}) = T(\vec{s}^{\,\bullet lm})$, i.e.,

$$T(\vec{s}^{\,lm}) = \frac{w^D_l}{\sum\limits_{k\in D} w^D_k} + \frac{\left(\sum\limits_{k\in D} w^D_k\right) - w^D_l - w^D_m}{\sum\limits_{k\in D} w^D_k}\,\frac{1}{1+q^{m-l}} = \frac{1}{1+q^{m-l}}$$

which implies $w^D_l q^{m-l} = w^D_m$. □

Note that the form of (3) implies that we do not have to translate a sequence so that its domain starts with one, in order to evaluate its rank function. This is because any extra powers of q in both the denominator and the numerator of the fraction $q^{k-1}/\sum_{l\in D(\vec{s})} q^{l-1}$ cancel out. One can now demonstrate the following corollary.

Corollary 2.5 *Let $\vec{s}_0, \vec{s}_1, \vec{s}_2$ be any three experience sequences such that $D(\vec{s}_0) \neq \emptyset$, $D(\vec{s}_0) \cap D(\vec{s}_1) = \emptyset$, and $D(\vec{s}_1) = D(\vec{s}_2)$. Then:*

(i) (Monotonicity.) $T(\vec{s}_0 \cup \vec{s}_1) < T(\vec{s}_0 \cup \vec{s}_2) \Leftrightarrow T(\vec{s}_1) < T(\vec{s}_2)$;

(ii) (Interpolation.) $T(\vec{s}_0) > T(\vec{s}_1) \Rightarrow T(\vec{s}_0) > T(\vec{s}_0 \cup \vec{s}_1) > T(\vec{s}_1)$.

(iii) (Amortization.) $|T(\vec{s}_0 \cup \vec{s}_1) - T(\vec{s}_0 \cup \vec{s}_2)| < |T(\vec{s}_1) - T(\vec{s}_2)|$. *Moreover, if we fix $\vec{s}_1$ and $\vec{s}_2$ and let $\|\vec{s}_0\| \to \infty$, then $T(\vec{s}_0 \cup \vec{s}_1) \to T(\vec{s}_0)$ and thus $|T(\vec{s}_0 \cup \vec{s}_1) - T(\vec{s}_0 \cup \vec{s}_2)| \to 0$.*

Proof. Let $\vec{s}_i = \{(k, e_i(k)) : k \in D(\vec{s}_i)\}$ for $i = 0, 1, 2$. Clearly, *(i)* follows directly from (3).

To prove *(ii)*, let $D = D(\vec{s}_0) \cup D(\vec{s}_1)$, and let $e(k) = e_0(k)$ if $k \in D(\vec{s}_0)$ and $e(k) = e_1(k)$ if $k \in D(\vec{s}_1)$; then by (3),

$$
\begin{aligned}
T(\vec{s}_0 \cup \vec{s}_1) = {} & \frac{\sum\limits_{i\in D(\vec{s}_0)} q^{i-1}}{\sum\limits_{j\in D} q^{j-1}} \sum_{k\in D(\vec{s}_0)} \frac{q^{k-1}}{\sum\limits_{i\in D(\vec{s}_0)} q^{i-1}}\, e(k) \\
& + \frac{\sum\limits_{i\in D(\vec{s}_1)} q^{i-1}}{\sum\limits_{j\in D} q^{j-1}} \sum_{k\in D(\vec{s}_1)} \frac{q^{k-1}}{\sum\limits_{i\in D(\vec{s}_1)} q^{i-1}}\, e(k)
\end{aligned}
$$

Setting $\alpha = \sum_{i\in D(\vec{s}_1)} q^{i-1}/\sum_{j\in D} q^{j-1}$ and $\beta = \sum_{i\in D(\vec{s}_2)} q^{i-1}/\sum_{j\in D} q^{j-1}$, we get $T(\vec{s}_1 \cup \vec{s}_2) = \alpha T(\vec{s}_1) + \beta T(\vec{s}_2)$, with $\alpha + \beta = 1$, which implies the claim.

To prove *(iii)*, as in Corollary 2.2, we use extension of T to sequences that can take negative values, and note that the sequence $(\vec{s}_0 \cup \vec{s}_1) - (\vec{s}_0 \cup \vec{s}_2)$, with domain $D = D(\vec{s}_0) \cup D(\vec{s}_1)$ is zero over $D(\vec{s}_0)$ and equal to $\vec{s}_1 - \vec{s}_2$ over $D(\vec{s}_1)$. Thus,

$$
\begin{aligned}
& T((\vec{s}_0 \cup \vec{s}_1) - (\vec{s}_0 \cup \vec{s}_2)) \\
= {} & \sum_{k\in D(\vec{s}_1)} \frac{q^{k-1}(e_1(k) - e_2(k))}{\sum_{i\in D} q^{i-1}} \\
= {} & \frac{\sum_{i\in D(\vec{s}_1)} q^{i-1}}{\sum_{j\in D} q^{j-1}} \sum_{k\in D(\vec{s}_1)} \frac{q^{k-1}(e^1(k) - e^2(k))}{\sum_{i\in D(\vec{s}_1)} q^{i-1}} \\
= {} & \frac{\sum_{i\in D(\vec{s}_1)} q^{i-1}}{\sum_{j\in D} q^{j-1}} \left(T(\vec{s}_1) - T(\vec{s}_2)\right)
\end{aligned}
$$

As $D(\vec{s}_0) \neq \emptyset$, we have $\sum_{i\in D(\vec{s}_1)} q^{i-1}/\sum_{j\in D} q^{j-1} < 1$ and thus $|T(\vec{s}_0 \cup \vec{s}_1) - T(\vec{s}_0 \cup \vec{s}_2)| < |T(\vec{s}_1) - T(\vec{s}_2)|$. Also, $\sum_{i\in D(\vec{s}_1)} q^{i-1}/\sum_{j\in D} q^{j-1} \to 0$ as $\|\vec{s}_0\| \to \infty$ and thus $T(\vec{s}_0 \cup \vec{s}_1) - T(\vec{s}_0 \cup \vec{s}_2) \to 0$. □

Note that *(i)* formalizes our intuition that the trust rank $T(\vec{s})$ should change monotonically with respect to addition of another sequence of experiences; *(ii)* formalizes our intuition that joining a sequence with a higher trust rank with a sequence of a lower trust rank should result in a sequence with a trust rank that is a degradation of the higher and an improvement of the lower. Finally, *(iii)* formalizes our intuition that difference in trust ranks of two experience sequences $\vec{s}_1$ and $\vec{s}_2$ "amortizes" when these sequences are extended with the same sequence $\vec{s}_0$; thus, initial few experiences have negligible impact on overall rank as the number of transactions increases.

If $q > 1$, then Corollary 2.5*(iii)* can be strengthened. In this case we can take only limited number of latest experiences to obtain an arbitrarily good approximation of $T(\vec{s})$. Let us denote by $\vec{s}[m, n]$ the restriction of $\vec{s}$ to the domain $D(\vec{s}) \cap \{m, \dots, n\}$.

Corollary 2.6 *If $q > 1$, then for every δ there exists K such that for every $\vec{s}$ satisfying $\|\vec{s}\| > K$, we have*

$$|T(\vec{s}) - T(\vec{s}[\,\|\vec{s}\| - K, \|\vec{s}\|])| < \delta.$$

Proof. For clarity, we fix the ranges of subscripts used in sums in this proof as follow:

- $m \in D(\vec{s})$;
- $k \in D(\vec{s}[1, \|\vec{s}\| - K - 1])$;
- $l, l_1 \in D(\vec{s}[\,\|\vec{s}\| - K, \|\vec{s}\|])$ and
- $i \in I_{\|\vec{s}\|-K-1}$.

Note that

$$\vec{s} = \vec{s}[1, \|\vec{s}\| - K - 1] \cup \vec{s}[\,\|\vec{s}\| - K, \|\vec{s}\|]$$

and

$$
\begin{aligned}
& \sum_l \frac{q^{l-1}\, e(l)}{\sum\limits_m q^{m-1}} \\
= {} & \frac{\sum\limits_{l_1} q^{l_1-1}}{\sum\limits_m q^{m-1}} \sum_l \frac{q^{l-1}\, e(l)}{\sum\limits_{l_1} q^{l_1-1}} \\
= {} & \frac{\sum\limits_l q^{l-1}}{\sum\limits_m q^{m-1}}\, T(\vec{s}[\|\vec{s}\| - K, \|\vec{s}\|]).
\end{aligned}
$$

Since

$$T(\vec{s}) = \sum_k \frac{q^{k-1}}{\sum\limits_m q^{m-1}}\, e(k) + \sum_l \frac{q^{l-1}}{\sum\limits_m q^{m-1}}\, e(l),$$

we have

$$\begin{aligned}
&T(\vec{s}) - T(\vec{s}[\|\vec{s}\| - N, \|\vec{s}\|]) \\
&= \sum_k \frac{q^{k-1}\, e(k)}{\sum\limits_m q^{m-1}} \\
&\qquad + \left(\frac{\sum\limits_l q^{l-1}}{\sum\limits_m q^{m-1}} - 1 \right) T(\vec{s}[\|\vec{s}\| - K, \|\vec{s}\|]) \\
&= \sum_k \frac{q^{k-1}\, e(k)}{\sum\limits_m q^{m-1}} - \left(\frac{\sum\limits_k q^{k-1}}{\sum\limits_m q^{m-1}} \right) T(\vec{s}[\|\vec{s}\| - K, \|\vec{s}\|])
\end{aligned}$$

Also, since $e(k) \le M$,

$$\sum_k \frac{q^{k-1}\, e(k)}{\sum\limits_m q^{m-1}} \le M \sum_k \frac{q^{k-1}}{\sum\limits_m q^{m-1}};$$

thus, by increasing the numerator and reducing the denominator, we get

$$\sum_k \frac{q^{k-1}}{\sum\limits_m q^{m-1}} < \sum_i \frac{q^{i-1}}{q^{\|\vec{s}\|-1}} < \frac{1}{(q-1)q^{K-1}}$$

Thus,

$$\begin{aligned}
&|T(\vec{s}) - T(\vec{s}[\|\vec{s}\| - K, \|\vec{s}\|])| \\
&\qquad \le 2M \frac{\sum\limits_k q^{k-1}}{\sum\limits_m q^{m-1}} \le \frac{2M}{(q-1)q^{K-1}}
\end{aligned}$$

which can be made arbitrarily small by taking sufficiently large K. □

Theorem 2.7 (Trust Update) *Let function $Q(D)$ be defined on finite sets of natural numbers so that $Q(D) = \sum_{i \in D} q^i$; let $n > max(D(\vec{s}))$ and $\vec{s}' = \vec{s} \cup \{(n, e(n))\}$. Then*

$$T(\vec{s}') = \frac{Q(D(\vec{s}))\, T(\vec{s})}{Q(D(\vec{s})) + q^{n-1}} + \frac{q^{n-1}\, e(n)}{Q(D(\vec{s})) + q^{n-1}}$$

Proof. By Theorem 2.4 we have

$$\begin{aligned}
T(\vec{s}') &= \sum_{k \in D(\vec{s})} \frac{q^{k-1} \vec{s}[k]}{\left(\sum\limits_{i \in D(\vec{s})} q^{i-1} \right) + q^{n-1}} \\
&\qquad + \frac{q^{n-1}\, e(n)}{\left(\sum\limits_{i \in D(\vec{s})} q^{i-1} \right) + q^{n-1}} \\
&= \sum_{k \in D} \frac{q^{k-1} \vec{s}[k]}{Q(D(\vec{s}))} \, \frac{Q(D(\vec{s}))}{Q(D(\vec{s})) + q^{n-1}} \\
&\qquad + \frac{q^{n-1}\, e(n)}{Q(D(\vec{s})) + q^{n-1}} \\
&= \frac{Q(D(\vec{s}))\, T(\vec{s})}{Q(D(\vec{s})) + q^{n-1}} + \frac{q^{n-1}\, e(n)}{Q(D(\vec{s})) + q^{n-1}}.
\end{aligned}$$

□

To obtain the new values, $T(\vec{s}')$ and $Q(D(\vec{s}'))$, we only need to keep track of two quantities, $T(\vec{s})$ and $Q(D(\vec{s}))$, rather than individual experience values $e_{ij}(k)$, for all $k \le n$. For $q > 1$

$$\begin{aligned}
&\frac{q^{n-1}}{S(D(\vec{s})) + q^{n-1}} \ge \frac{q^{n-1}}{\sum_{k=1}^{n-1} q^{k-1} + q^{n-1}} \\
&\qquad = 1 - \frac{q^{n-1} - 1}{q^n - 1} = 1 - \frac{1 - 1/q^{n-1}}{q - 1/q^{n-1}} > 1 - \frac{1}{q}
\end{aligned}$$

and consequently $Q(D(\vec{s}))/(Q(D(\vec{s})) + q^{n-1}) < 1/q$. Thus, as q increases, the dependency of the trust rank on most recent samples also increases, i.e., the value of q controls how fast the trust rank updates. The proper value for q should be chosen close to but larger than one, on the basis of statistical data from the community of agents involved, matching the expected volatility of agents and striking the right balance between importance of longer term performance and likelihood of sudden changes in agent's behavior.

2.3 Asymptotic Behavior of Trust Rank

Introducing trust rank function for infinite experience sequences allows an analysis of the asymptotic behavior of trust rank functions on finite sequences. Thus, let $\vec{S}$ be a function from an infinite domain $D(\vec{S}) \subseteq \mathbb{N}$ with values in $[0, M]$. We denote by $\vec{S}[m, n]$ the restriction of $\vec{S}$ to the domain $D(\vec{S}) \cap \{m, \ldots, n\}$. Let $\vec{S}$ be any infinite experience sequence with the domain $D(\vec{S})$; we define $T(\vec{S}) = \lim_{n \to \infty} T(\vec{S}[1, n])$ whenever such limit exists. Assume $n \in \mathbb{N}$; then $\vec{S}\downarrow_n$ is the infinite sequence with the domain $D(\vec{S}\downarrow_n) = \{k : (k + n) \in D(\vec{S})\}$ defined by $\vec{S}\downarrow_n [k] = \vec{S}[k + n]$.

Theorem 2.8 *For every integer n, $T(\vec{S})$ is defined if and only if $T(\vec{S}\downarrow_n)$ is defined, in which case $T(\vec{S}) = T(\vec{S}\downarrow_n)$.*

Theorem 2.8 asserts that $T(\vec{S})$ is shift invariant not only for shifts in the positive direction but also in the negative direction, when initial few terms could be dropped. It can also be easily seen that $T(\vec{S})$ for infinite sequences also satisfies the Consistency axiom and the following strengthening of the Averaging Axiom: $\liminf(\vec{S}) \leq T(\vec{S}) \leq \limsup(\vec{S})$.

Theorem 2.9 *If experiences in an infinite sequence of experiences $\vec{S}$ converge, i.e., if $\lim_{n\to\infty} \vec{S}[n] = s$, then $\lim_{n\to\infty} T(\vec{S}[1,n])$ also exists and equals s.*

Proof. Let $\delta > 0$ be arbitrary, and let $K \in \mathbb{N}$ be such that $|\vec{S}[n] - s| < \delta$ for all $n > K$. Using Corollary 2.5 *(iii)* with $\vec{s}_0 = \vec{S}[K+1,n]$ and $\vec{s}_1 = \vec{S}[1,K]$ we can chose $L > K$ such that for all $n > L$, $|T(\vec{S}[1,n]) - T(\vec{S}[K,n])| < \delta$. Let $\vec{s}^{\,n}$ be such that $D(\vec{s}^{\,n}) = D^n = D(\vec{S}) \cap \{K,\ldots,n\}$, and $\vec{s}[k] = s$ for all $k \in D^n$; then $\vec{s}^{\,n}[k] - \vec{S}[K,n][k] < \delta$ and thus $|T(\vec{s}^{\,n}) - T(\vec{S}[K,n])| = |T(\vec{s}^{\,n} - \vec{S}[K,n][k])| < \delta$. Hence, $|T(\vec{S}[1,n]) - T(\vec{s}^{\,n})| < 2\delta$, i.e., $|T(\vec{S}[1,n]) - s| < 2\delta$, which implies our claim. $\square$

Theorem 2.9 states that $T(\vec{S})$ exists for every convergent infinite experience sequence and equals to the limit of the sequence. The converse is false; consider $\vec{S}^{\bullet}$ with domain $\mathbb{N}$ such that for $\vec{S}^{\bullet}[2k+1] = 1$ and $\vec{S}^{\bullet}[2k] = 0$ all k. It is easy to see that $T(\vec{S}^{\bullet})$ converges to $1/2$ for $q = 1$. Thus, the notion of reputation rank is a form of a generalized infinite impulse response averaging filter applicable to "partial signals" which need not be defined for all instants in time.

2.4 Weight of Pairwise Evidence

Having a high trust rank on the basis of just a few transactions long time ago is clearly not the same as having a high reputation rank on the basis of a large number of recent transactions. To formalize this intuition, we introduce the notion of the *weight of pairwise evidence an agent A_i has about the trustworthiness of another agent A_j up to the instant n*, denoted by $w_{ij}(n)$. Thus, let $\vec{s}_{ij}(n)$ be the experience sequence of A_i with A_j that takes into account **all** transactions A_i has had with A_j up to the moment n. Then

$$w_{ij}(n) = \sum_{k \in D(\vec{s})} q^{k-n},$$

where q is the same constant as in (2.4). In essence, we count the number of transactions of A_i with A_j, discounting each transaction by the factor $1/q^{n-k}$ where $k \in D(\vec{s}_{ij}(n))$ is the instant of transaction. In this way, the larger the value of $w_{ij}(n)$ the more significant and reliable the value $T(\vec{s}_{ij})$ is at an instant n. Note that w_{ij} satisfies the recursion:

$$w_{ij}(n+1) = \begin{cases} q\,w_{ij}(n) + 1 & \text{if } n+1 \in D(\vec{s}_{ij}(n+1)) \\ q\,w_{ij}(n) & \text{otherwise} \end{cases}$$

Thus, $w_{ij}(n)$ can also be evaluated recursively; for $q = 1$ we have $w_{ij}(n) \leq n$, and for $q > 1$, $w_{ij}(n)$ are bounded because $w_{ij}(n) \leq \sum_{k=0}^{\|\vec{s}_{ij}(n)\|-1} q^{-k} \leq \sum_{k=0}^{\infty} q^{-k} = q/(q-1)$.

3 Aggregation in Community

We can now tackle the problem of averaging in community. Let $\vec{s}_{ij}(n)$ be the sequence of experiences in *all* transactions of reporting agent A_i with agent A_j up to and including instant of time n. We assume that the agency has calculated the trust rank $\tau_{ij}(n) = T(\vec{s}_{ij}(n))$ of each agent A_j with every other agent A_i which were involved in at least one transaction up to the instant n, and the corresponding weight of pairwise evidence $w_{ij}(n)$. The agency now has to assign a *community reputation rank* $1 \leq \rho_j(n) \leq M$ to each agent A_j that had transactions up to the instant n, as well as a *weight of community evidence $W_j(n)$ for the reputation of A_j* on the basis of all available values $\tau_{ij}(n)$ and $w_{ij}(n)$.

Let $\Delta_j(n)$ denote the set of indices of all agents with which agent A_j has had transactions up to the instant n. The community reputation rank should be a form of an average over the community of agents in the sense that $\min\{\tau_{ij}(n) : i \in \Delta_j\} \leq \rho_j(n) \leq \max\{\tau_{ij}(n) : i \in \Delta_j\}$. The weight of community evidence $W_j(n)$ for the reputation of A_j should be a cumulative weight of pairwise evidence obtained by aggregating all weights $\{w_{ij}(n) : i \in \Delta_j\}$ and that reflects the confidence we can have in accepting the value of the reputation rank $\rho_j(n)$.

So in our setup we can assume that at instant n the agency has a partial table $\mathcal{T}_n$, such that for each $i, j \leq n$ the (i,j)-entry of the table $\mathcal{T}_k$ is either blank if A_i and A_j were not involved in a transaction up to the instant n, or is a pair $(\tau_{ij}(n), w_{ij}(n))$, $\ i \neq j$. We now need a mapping $\boldsymbol{\rho}(\mathcal{T}_n) \mapsto \mathcal{R}_n$ assigning to each partial table $\mathcal{T}_n$ a partial table $\mathcal{R}_n$ of pairs $(\rho_i(n), W_i(n))$ assigned for instant n. The table $\mathcal{R}_n$ is made available to all agents to help them estimate the risks involved in entering into a transaction with other agents.

To explain the essence of our method of arriving at the reputation rank $\rho_j(n)$ and the cumulative weight of community evidence $W_j(n)$, we consider the following situation. Assume that we have a community of agents $\mathcal{A} = \{A_i : i \leq N\}$ with already assigned ranks $\{\rho_i(n) : i \leq N\}$, and an additional agent A_0 with unknown rank. Assume that A_0 has had transactions with all of the agents in $\mathcal{A}$, and that they have submitted their experience reports $e_{i0}(n)$. Thus, we can assume that we have calculated the trust ranks $\tau_{i0}(n)$

of A_0 with every agent $A_i \in \mathcal{A}$, and the weight of evidence $w_{i0}(n)$ for such trust ranks. Then the reputation rank $\rho_0(n)$ of A_0 can be evaluated using the following weighted average:

$$\rho_0(n) = \sum_{i=1}^{N} \frac{w_{i0}(n)\rho_i(n)\ \tau_{i0}(n)}{\sum_{k=1}^{N} w_{k0}(n)\rho_k(n)}. \tag{4}$$

We see that individual trust ranks $\tau_{i0}(n)$ of agent A_0 with agents A_i, $i \leq n$, are averaged with weights that reflect reputation $\rho_i(n)$ of the reporting agents, as well as the amount of evidence they have to support the values $\tau_{i0}(n)$. In this way, the highest impact on the value of the community reputation $\rho_0(n)$ of agent A_0 will have agents with high community reputation rank $\rho_i(n)$ at the given instant n, which also have the large weight of evidence $w_{i0}(n)$ backing their corresponding estimate τ_{i0} of trustworthiness of agent A_0. One can also define the cumulative, community weight of evidence for the reputation $\rho_0(n)$ of agent A_0 as

$$W_0(n) = \sum_{k=1}^{N} w_{k0}(n)\rho_k(n) \tag{5}$$

as a measure of weight of total evidence that supports rank $\rho_i(n)$, because each individual weight $w_{k0}(n)$ is prorated by the reputation of the reporting agent. Thus, (4) can be written as

$$\rho_0(n) = \sum_{i=1}^{N} \frac{w_{i0}(n)\rho_i(n)\ \tau_{i0}(n)}{W_0(n)} \tag{6}$$

and interpreted that the reputation rank $\rho_0(n)$ is a weighted sum of trust ranks τ_{i0} of agent A_0 with agents A_i, $i \leq N$, with weight given to an agent A_i that reflects its share of prorated evidence $w_{i0}(n)\rho_i(n)$ in the total of prorated evidence $W_0(n)$.

One may object that reputation ranks $\rho_i(n)$ and weights of evidence $w_{i0}(n)$ are not on a par; for example, one might argue that halving the reputation rank of a reporting agent A_i should have more severe impact than halving the weight of its evidence $w_{i0}(n)$. Thus, we generalize the (4) by allowing an arbitrary continuous scaling function $f(w,\rho) : [0, q/(q-1)] \times [1, M] \rightarrow \mathbb{R}^+$ that is strictly increasing in both arguments and satisfies $w, \rho \leq f(w,\rho)$ and by letting

$$\rho_0(n) = \sum_{i=1}^{N} \frac{f(w_{i0}(n), \rho_i(n))\ \tau_{k0}(n)}{\sum_{k=1}^{N} f(w_{k0}(n), \rho_k(n))}. \tag{7}$$

For example, in our numerical experiments $f(w,\rho) = w^\alpha \rho^\beta$ performed very well, with α, β depending on particular value of the discounting factor q appearing in our Representation Theorem 2.4.

Note that the above does not solve our problem of assigning the reputation ranks $\rho_i(n)$, because we assumed that we had such ranks for all agents except A_0, and we did not take into account at all experiences of A_0 with other agents. To solve this problem we look for a solution of the following system of equations in variables $\boldsymbol{\rho}_i$, $i \leq n$.

$$\left\{ \sum_{i \in \Delta_j(n)} \frac{f(w_{ij}(n), \boldsymbol{\rho}_i)\ \tau_{kj}(n)}{\sum_{k \in \Delta_j(n)} f(w_{kj}(n), \rho_k)} = \boldsymbol{\rho}_j \right\}_{j \leq N} \tag{8}$$

Thus, if $\rho_i(n)$, $i \leq N$ are chosen to be the solutions of the system (8), then for *every* agent A_j its community reputation rank $\rho_j(n)$ at an instant n is equal to the weighted average of all trust ranks $\tau_{ij}(n)$ of A_j at instant n, with every other agent A_i that was involved in transactions with A_j up to the moment n, weighted by A_i's share of contribution to the total community evidence, i.e., $f(w_{ij}(n), \rho_i)/\sum_{k \in \Delta_j(n)} f(w_{kj}(n), \rho_k)$, and which is commensurate to both its A_i's reputation rank $\rho_i(n)$ and its weight of evidence $w_{ij}(n)$. We now make an inessential change to possible range of our experience estimates $e_{ij}(n)$, allowing only values $1 \leq e_{ij}(n) \leq M$, with $M >> 1$; for example, we can set $1 \leq e_{ij}(n) \leq 100$. Then, by our Averaging axiom, also the values $\tau_{ij}(n)$ are in the same range. Consider now the mapping F given by

$$F : (\rho_i \ : \ i \leq N) \mapsto \left(\sum_{i \in \Delta_j(n)} \frac{f(w_{ij}(n), \rho_i)\ \tau_{kj}(n)}{\sum_{k \in \Delta_j(n)} f(w_{kj}(n), \rho_k)} \ : \ j \leq N \right).$$

Since for every j the weights in the weighted j^{th} sum add up to one and since the scaling function f is continuous in both variables, it is easy to see that F is a continuous mapping of the N-dimensional cube $[1, M]^N$ into itself. Thus, since F is a continuous map from a convex and compact subset of R^N into itself, by the Brouwer Fixed point theorem F must have a fixed point, see e.g., [3], which we take as the values of $(\rho_j(n) \ : \ j \leq N)$. The fixed point can be obtained by standard iterative procedures, and our experimental results show that the method is both fast and numerically robust. It is possible to reset the clock without causing disruption and to partition the market into overlapping sub communities to prevent an explosion in size of the system of equations. The details, together with the issues of software engineering of such a system will be presented separately in a forthcoming paper.

Finally, we can evaluate the community evidence $(W_j(n) \ : \ j \leq N)$ using

$$W_j(n) = \sum_{k \in \Delta_j(n)} f(w_{kj}(n), \rho_k);\ j \leq N. \tag{9}$$

Thus, at each moment n and for each $j \leq N$, $\rho_j(n)$ is a representation and an estimate for our informal notion of community reputation enjoyed by agent A_j, and $W_j(n)$ is a community based measure of reliability of such reputation estimate.

One might ask a question how one should use reputation rank and the weight of evidence to make a decision. For example, given two agents, one with reputation of 50 and weight of evidence $.5$ and another with reputation rank 60 and weight of evidence $.4$, which should one chose? The answer depends on the particular features of the community of agents, and must be probabilistic in nature, in the sense that agents will act so that they maximize *the expected* utility in transactions. These are problems that merit a separate paper.

Conclusion. Tennenholtz's [7] idea on the reputation of agents can be paraphrased as follows. Say an agent A_j is *supported* by a set S_j of other agents if the agents in S_j provide high trust rankings for A_j. Since his discussion is about a static domain there is no notion of time progression. Moreover, he argued that support by agents which themselves have high reputations should result in higher reputation for the supported agent than support by low reputation agents. We feel that we have provided an adequate and rigorous formalization of this idea, also extending Tennenholtz's static setup to one that involves time sequences of transactions.

References

[1] R. Falcone and C. Castelfranchi. Social trust: a cognitive approach. In C. Castelfranchi and Y.-H. Tan, editors, *Trust and deception in virtual societies*, chapter 3, pages 55–90. Kluwer Academic Publishers, 2001.

[2] N. Foo and J. Renz. Experience and trust: A systems-theoretic approach. In *18th European Conference in Artificial Intelligence*, 2008.

[3] D. H. Griffel. *Applied Functional Analysis*. Dover Publications, Jun 2002.

[4] C. M. Jonker and J. Treur. Formal analysis of models for the dynamics of trust based on experiences. In F. J. Garijo and M. Boman, editors, *Multi-Agent System Engineering: Proceedings of the 9th European Workshop on Modelling Autonomous Agents in a Multi-Agent World, MAAMAW'99, Valencia, Spain, June 30 - July 2 1999*, volume 1647 of *Lecture Notes in Artificial Intelligence*, pages 221–232. Springer-Verlag, 1999.

[5] A. Schlosser, M. Voss, and L. Brückner. On the simulation of global reputation systems. *Journal of Artificial Societies and Social Simulation*, 9(1):4, 2005.

[6] G. C. Silaghi, A. E. Arenas, and L. M. Silva. Reputation-based trust management systems and their applicability to grids. Technical Report TR-0064, CoreGRID, Feb 2007.

[7] M. Tennenholtz. Reputation systems: an axiomatic approach. In *AUAI '04: Proceedings of the 20th conference on Uncertainty in artificial intelligence*, pages 544–551, Arlington, Virginia, United States, 2004. AUAI Press.