

Integrating Indicators of Trustworthiness into Reputation-Based Trust Models Insurance, Certification, and Coalitions

Sascha Hauke, Florian Volk, Sheikh Mahbub Habib, and Max Mühlhäuser

Technische Universität Darmstadt/CASED,
Telekooperation, Hochschulstraße 10, 64283 Darmstadt
{sascha.hauke,florian.volk,sheikh.habib}@cased.de,
max@informatik.tu-darmstadt.de

Abstract. Reputation-based trust models are essentially reinforcement learning mechanisms reliant on feedback. As such, they face a cold start problem when attempting to assess an unknown service partner. State-of-the-art models address this by incorporating dispositional knowledge, the derivation of which is not described regularly. We propose three mechanisms for integrating knowledge readily available in cyber-physical services (e.g., online ordering) to determine the trust disposition of consumers towards unknown services (and their providers). These reputation-building indicators of trustworthiness can serve as cues for trust-based decision making in eCommerce scenarios and drive the evolution of reputation-based trust models towards trust management systems.

1 Introduction

Internet-based and mediated services have managed to capture considerable market shares in what used to be primarily real-world markets. The further amalgamation of online and real-world service provisioning, such as online ordering of physical goods, e.g., books, or provisioning of services, be they hotel bookings or cloud compute services, promise additional convenience for consumers and business opportunities for providers. Personal and institutional procedures for evaluating whom to trust in this new environment are still in the process of being established. The relative ease of setting up an online business, as compared to brick-and-mortar enterprises, leads to more transience in a market.

In order to overcome these challenges and build trust in unregulated online markets, such as the present and future internet, two distinct schools of thought have emerged. On the one hand, the “hard” approach to trust dictates rigorous certification and provable chains of credentials between a (presumably) entirely trusted root and a node. This is used, for instance, in trusted computing applications. On the other hand, the “soft” way of thinking about trust relegates trust to the domain of probabilities, conventionally stating that trust is a *subjective probability* [4] of somebody else acting as expected. This probability is typically derived from feedback histories using (probabilistic) trust models, such as [11].

In its current form, neither is entirely satisfactory when addressing the needs of (future) internet-based markets. While hard trust might be sufficient to provide information on the identity of another entity, possibly its persistence and even some of its capabilities, its shortcomings are in describing the behavior of that other entity. Soft trust, with its reliance on feedback and reputation, expressed as community standing, is prone to particular attacks. It also faces shortcomings such as those related to reinforcement learning.

In this paper, we present an extension to the established CertainTrust trust model [22]. The concepts of insuring, certifying, and coalition forming are adapted to be used as an extension to the model. By explicitly modeling cues that are already well-established in real-world interactions for use in a reputation-based trust model, the approach contributes to mitigating the cold start/market entry problem. Additionally, by allowing providers to represent their trustworthiness, the modeling of these approaches forms a first step of evolving CertainTrust into a trust management system (following the definition of such a system by Jøsang et al. [12]). By integrating certification processes with reputation-based trust, an integration of hard and soft trust approaches is potentially enabled.

The impacts on trust and reliance are discussed in the context of a cyber-physical service provision context. They are, furthermore, briefly presented in a qualitative agent-based simulation. Insurance and certification models were chosen, because for both there exist functioning real-world markets with highly-reputable service providers. These providers can serve as persistent trust anchors for more transient online services, such as cloud-based offerings by small and medium enterprises.

The contribution can, of course, be adapted to other reputation-based trust models and is not limited to the given use case by any means.

The remaining document is structured as follows: section 2 presents a use case for the proposed approach presented in section 3. Section 4 discusses the application of the indicators to the use case and presents the results of an agent-based simulation qualitatively showing the effect of the individual operators. Section 5 surveys related work. In section 6, some conclusions are drawn.

2 Use Case

For the use case, consider a customer trying to establish trust on a cyber-physical service. Furthermore, suppose that the customer does not have any prior experience with that particular service. It is therefore not immediately possible to derive the trustworthiness of the service provider from direct experience. In order to derive the reliability of the service, the conventional approach for reputation-based trust models (cf. e.g., [9,11,22]) is to query trusted witnesses for information. However, even in the absence of reliable witnesses, both initial reliability and decision trust [13] can be established from other cues.

In cyber-physical services, that involve both digital and real-world processes, such as online ordering and physical shipping of goods, service delivery is generally not monolithic. Rather, the service provisioning processes can be sub-divided

into sub-components, some of which are visible to the customer and may be associated with distinct entities on which trust can be established individually.

Figure 1 outlines a general scenario in which a customer establishes trust on an unknown (foreign) composite service. By necessity, several components of the service are visible to the customer, such as payment/billing and shipping agents used by the service provider. We assume that the billing process is handled through an intermediary, specifically a credit card company. For the core service provisioning process, we further assume that the composite service provider chooses not to reveal its internal processes to the customer directly. It may, however, use an external auditing and certification provider (e.g., ISO) to certify its internal processes. In this paper, we abstract from the multi-dimensionality of trust. Thus, a certification is considered to be representative for the reliability of the internal service provisioning process.

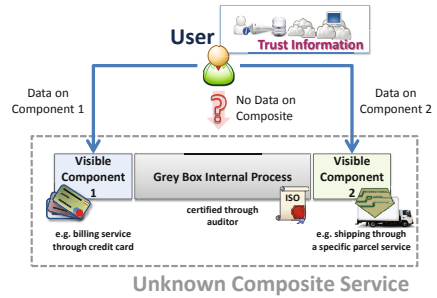


Fig. 1. Use Case: Cyber-Physical Service Composition

3 Approach

Meanings and definitions of trust have been discussed at some length in the literature (cf. e.g., [4,19,18]). Within the scope of this paper, we will follow [12] in differentiating reliability trust and decision trust. We will define reliability trust according to Gambetta [4,12]:

Definition 1. *Trust is the subjective probability by which an individual expects that another individual performs a given action on which its welfare depends.*

In particular, we consider trust to be an adequate approximator of trustworthiness. The expectation value E computed by the CertainTrust trust model represents such a trust score. When having to make a decision, however, further considerations are involved, beyond the supposed reliability expressed by the trust score. This is reflected in decision trust [12]:

Definition 2. *Trust is the extent to which a given party is willing to depend on something or somebody in a given situation with a feeling of relative security, even though negative consequences are possible.*

Reliability trust can be said to *inform* decision trust. However, risk, gain, loss and reliance [20] are also contributing to the decision-making process. Consequently, decision trust will be modeled using expected utility theory [13,17]. The probabilities, denoted as p , used in the computation of the expected utility will be derived from reliability trust. In particular, the values of various instances of p , e.g., used in equations 2 and 4, are approximated by the reliability trust score from CertainTrust.

Let G be a benefit expected from an interaction, i.e., the positive gain, and L the corresponding loss, or negative gain. Furthermore, let $p \in [0, 1]$ be the probability of a beneficial outcome. Then, the expected utility EU of an interaction can be defined as [13,17]:

$$EU := p \cdot G - (1 - p) \cdot L \quad (1)$$

3.1 Using CertainTrust to Measure Reliability Trust

To model experiences as trust information, the CertainTrust model and the Human Trust Interface (HTI) from [22] are applied. CertainTrust models trust as opinions based on positive evidences r_i and negative evidences s_i . Using collected evidence (e.g., feedback), it allows to calculate an expectation value $E_{f,w,N}(t, c) = t \cdot c + (1 - c) \cdot f$. The certainty c depicts on how much evidence the trust value t is based. A low amount of evidence (low certainty c) is compensated by using the (dispositional) initial trust value f . The parameter w allows to express the weight of dispositional trust, while N denotes the maximal amount of expected evidence, in this paper’s case: the amount of single experiences.

The true value of the probability p can be considered an inherent quality of an entity that cannot be measured directly. It is assumed that $E_{f,w,N}$ is an appropriate approximator for p . In the following, various variables – e.g., c_{issuer} , t_{issuer} , $c_{candidate}$, $t_{candidate}$, and f – are derived using CertainTrust. In particular, they do not have to be determined manually.

3.2 Using Expected Utility to Model Decision Trust

Consumers selecting a service will generally try to maximize their utility. Thus, they will tend to select the service with the highest expected utility EU . The expected utility function is subject to uncertainty, because $E_{f,w,N}$ is used instead of the true value for p . Most variables here are either direct results of applying the trust model, are derived through the delegation mechanisms discussed in the following or are explicitly available from the context of an interaction (e.g., the premium a service provider charges for the use of a credit card, which offers an insurance option, would cover $L_{insurer}^{fix}$).

Prior experiences by consumers and indicators of trustworthiness are bound to service providers’ identities. Therefore, persistent identities are desirable. Otherwise, bad reputation could easily be “whitewashed” by re-entering the market with a new identity [2]. An upfront monetary investment bound to an identity shows the dedication of a service provider to this identity and therefore reflects an incentive to act trustworthy [3]. Unlike the basic approach [3], that requires a trusted third party or a managed marketplace to bind an investment to an identity, our approach solely relies on “trust-building” services, e.g., insurance services and certification services.

3.3 Reliance through Insurance

The insurance case relies on three entities: The *consumer* trying to identify the most appropriate service provider to select, the *service provider* under evaluation,

and an *insurance provider* insuring the transaction if the consumer decides to interact with the service provider. The relations between the entities are outlined in figure 2. Insurance provides reliance [20], and thus affects decision trust, by reducing the risk of asset loss attendant with an interaction. It therefore should contribute to “[...] a feeling of relative security [...]” (cf. definition 2).

Let $p_{candidate}$ be the probability of a successful interaction with a candidate service provider, and $p_{insurer}$ the probability of a successful interaction with an insurance provider that vouches or guarantees the interaction between consumer (acting as the initiator [22]) and the service provider (acting as the candidate). Furthermore, let the cost, or negative gain, the consumer experiences in case of an unsuccessful interaction with the service provider, be denoted $L_{candidate}$.

Analogously, $L_{insurer}^{fix}$ is the cost (if any) of the insurance contract to the consumer. Additionally, $L_{insurer}^{var}$ indicates the expenses incurred by the consumer when making an insurance claim against a failed interaction. In this case, the expected utility of the interaction for the consumer is:

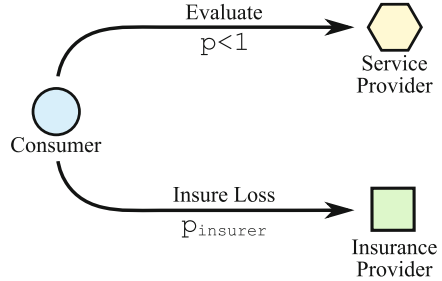


Fig. 2. Trust Delegation with Insurance

$$\begin{aligned}
 EU &:= p_{candidate} \cdot G \\
 &- (1 - p_{candidate})(1 - p_{insurer}) \cdot (L_{candidate} + L_{insurer}^{var}) \\
 &- (1 - p_{candidate})(p_{insurer}) \cdot L_{insurer}^{var} \\
 &- L_{insurer}^{fix}
 \end{aligned} \tag{2}$$

Table 1. Reputation Updates with Insurance

Interaction		Update	
Provider	Insurer	Provider	Insurer
success	-	positive	-
failure	success	negative	positive
failure	failure	negative	negative

After an insured interaction between a consumer and the selected candidate took place, the consumer updates its trust values according to table 1. In case the interaction with the provider succeeded, additional positive evidence regarding the provider is created, e.g., by increasing the value of $r_{provider}$ by 1. In this successful case, action from the insurer is not demanded and no further evidence regarding the insurer is collected. However, if the interaction with the selected candidate fails, there are two possible cases. If the insurer is called upon and reimburses $L_{candidate}$ to the consumer, therefore compensating the negative gain for the consumer, new positive evidence for the insurer is collected. If the insurer

fails in compensating the negative gain, new negative evidence regarding the insurer is collected, e.g., by increasing the value of $s_{insurer}$. In both cases, new negative evidence regarding the selected provider is created analogously.

3.4 Assessing Reliability through Certification

Similar to the insurance case from the previous section, this case consists of three interacting entities. The consumer is evaluating a service provider for selection. This service provider is certified by a certification provider the consumer has prior knowledge about but does not interact directly with (see figure 3).

For this paper, we assume a certification provider certifies service quality for an entire service or service component. We abstract from the multi-dimensionality of trust at this point. Certification of partial aspects of a service (component) can be combined into an overall rating, for instance using the propositional logic operators of CertainLogic [23]. Formally, a certification describes a specific minimum level of quality as $q_{cert} \in [0, 1]$ that a certification provider awards to the certified party, ideally after completing an audit.

This kind of limited trust delegation, employing a “probabilistic” certificate value and a certification provider that is not necessarily a completely trusted third party, influences the reliability trust for the candidate. In particular, in order to preserve the importance of direct experience over other kinds of information, we propose to include certification information in the initial expectation value f of CertainTrust. In its simplest form, it thus follows:




Fig. 3. Trust Delegation with Certification

$$\begin{aligned}
 p_{issuer} &= E(t_{issuer}, c_{issuer}) \\
 &= c_{issuer} \cdot t_{issuer} + (1 - c_{issuer}) \cdot f \\
 f_{cert} &= \max(f, \min(p_{issuer}, q_{cert}))
 \end{aligned} \tag{3}$$

$$E^{cert}(t_{candidate}, c_{candidate}) = c_{candidate} \cdot t_{candidate} + (1 - c_{candidate}) \cdot f_{cert}$$

The variables c_{issuer} , t_{issuer} , $c_{candidate}$, $t_{candidate}$, and f are derived using CertainTrust. In particular, they do not have to be determined manually.

The modified reliability trust score $E^{cert}(t_{candidate}, c_{candidate})$ informs the decision trust. Let $p_{candidate}^{cert} = E^{cert}(t_{candidate}, c_{candidate})$ be the probability of a successful interaction with a candidate service provider, given a certification from a certification provider. Then, the expected utility of the interaction between a consumer and a certified service provider can simply be described as:

$$EU := p_{candidate}^{cert} \cdot G - (1 - p_{candidate}^{cert}) \cdot L \tag{4}$$

Table 2. Reputation Updates with Certification

Interaction		Update	
Provider	Certifier	Provider	Certifier
success	–	positive	positive
failure	–	negative	negative

Trust evidence updates after an interaction (as per [22]) are created according to table 2, taking into account only the performance of the selected provider. However, new trust evidence is created for both the selected provider and the certifier. Thus, while trust is delegated from the certification provider to the candidate service provider, trust updates are delegated from the service provider to the certification provider. In case of a negative outcome, the new evidence regarding the certifier is justified because the certification was incorrect for at least this interaction. While being unable to determine if this incorrect certification holds for all cases, it is perceived by the consumer as an incorrect certification for the selected provider. Thus, the certifier might also fail to certify other providers correctly, e.g., due to shortcomings in the certification or auditing process.

3.5 Joint Reliability through Coalitions

Another way for service providers to represent their trustworthiness is the formation of coalitions with other service providers. The motivation behind the introduction of this mechanism is the underlying assumption that a mutual association with another trustworthy provider serves as an indicator of trustworthiness. Lack of experience with one service provider, i.e., the candidate, can thus be compensated by the consumer, i.e., the initiator, via the delegation of trust from associated service providers, i.e., its associates, that might be known to the consumer.

While a coalition is different from an upfront monetary investment as insurance or certification, it is unlikely that established providers form coalitions with service providers that are unknown to them. Sybil attacks from malicious service providers that spawn many identities and create coalitions between them are unlikely – because they are ineffective: coalitions influence the probability of being selected by increasing the visibility of a service provider. Being associated with a *well-known and trusted* party becomes an implicit certification. A mutual coalition of unknown service providers does not increase the visibility of the participants.

Assume a consumer wishes to evaluate a candidate service provider. It lacks, however, past direct experiences and recommendations to form a reliable opinion. This lack of knowledge might lead the consumer to choose another, better known service provider or forgo the interaction altogether. In order to alleviate the problem and be able to realize a profit from the interaction, it is in the candidate's best interest to increase the consumer's perception of its trustworthiness. To this end, the candidate presents a list of other service providers it is associated with in a coalition to the consumer. As shown in figure 4, this is done under the expectation that the consumer has prior experiences with at

least some of those. In this case, the experience the consumer has in the service provider's associates is transferred to the candidate.

Realizing Mutual Coalition. In composed services, coalitions are already in place. By taking into account the nature of the cooperation of service composition sub-components and their respective providers, trust delegation through the proposed coalition mechanism is a feasible method of establishing trust. Whether or not such a delegation is appropriate is dependent on the direction of the trust delegation with regard to the order of the sub-components within the process, as well as on power symmetries and enforcement possibilities among the providers associated within a service composition. For instance, considering the use case in section 2, it can be argued that the credit card provider (i.e., visible component 1 in figure 1) is strongly connected to the grey box internal process. This is due to strong obligations and enforcement mechanisms (e.g., binding legal agreements and litigation possibilities) integrating the respective service providers.

If not explicitly cooperating in the service composition under evaluation, service providers that otherwise cooperate can enable coalition-based trust delegation through the following mechanism by advertising their cooperation to the customer. The customer, acting as initiator, can consequently verify the coalitions and transfer trust accordingly.

Mutual coalitions are realized through the exchange and mutual acknowledgment of cooperation messages. A process for this is depicted in figure 5.

1. Service provider A creates a message $m_{A,B} = \langle UID_A, UID_B, data \rangle$ consisting of
 - a unique identifier representing provider A , e.g., an X.509 certificate
 - a unique identifier representing associate B , e.g., an X.509 certificate.
2. Service provider A forwards $m_{A,B}$ to service provider B .
3. B acknowledges its coalition with A by signing $m_{A,B}$.
4. B returns the signed cooperation message $\{m_{A,B}\}_{sigB}$.
5. A forwards its signed counterpart cooperation message $\{m_{A,B}\}_{sigA}$.

These cooperation messages can then be presented to potential consumers, in order to facilitate the coalition-based trust delegation.

6. A potential consumer C evaluating service provider A requests indicators of trustworthiness from A .
7. A supplies C with a list of cooperation messages.
8. C may validate the coalition between A and B by requesting B to verify the signed cooperation message $\{m_{A,B}\}_{sigB}$.




Fig. 4. Trust Delegation with Associates

9. Service provider B , as an associate of A , either confirms or denies the coalition with A , in particular regarding both the validity of the signature and currentness of the coalition.
10. The consumer C delegates the trustworthiness of B to A .

Delegating Trust in Coalitions.

Let $E_{f,w,N}(t_{candidate}, c_{candidate}) \approx f$ with certainty $c_{candidate} \approx 0$ be an estimate for $p_{candidate}$. $f \in [0, 1]$ represents the initial trust disposition of the consumer [22], which is conventionally chosen conservatively low. Thus, for a trustworthy candidate, it should typically hold that if $c_{candidate} \rightarrow 1$, then $p_{candidate} \rightarrow t_{candidate} \gg f$. $t_{candidate}$ is the average of prior experiences the consumer had with the candidate, each of which can be either positive or negative. Let $r_{candidate}$ and $s_{candidate}$ be the sum of positive and negative experiences, respectively [22]. Then, $t_{candidate} = \frac{r}{r+s}$.

The condition that $c_{candidate} \approx 0$ implies that $r_{candidate} + s_{candidate} \ll N$, where N is a constant denoting the minimum number of experiences required to reach a certainty $c_{candidate}$ of 1, as per [22]. In the proposed coalition scheme, the gap between $r_{candidate} + s_{candidate}$ and N is to be filled with experiences on associated service providers.

Let associates A_1, \dots, A_m be service providers associated with the candidate provider. Furthermore, let $(r_{A_i}, s_{A_i}), i \in 1, 2, \dots, m$ be the positive and negative experiences the consumer has made with service provider A_i . In order to minimize inequality effects regarding the number of experiences that influence trust delegation, we apply a normalization in the same manner as [22]:

$$norm_N(r, s) = \begin{cases} 1 & \text{if } r + s \leq N \\ \frac{N}{r+s} & \text{else} \end{cases}$$

$$\tilde{r} = r_{candidate} + \delta \cdot \alpha \sum_{i=1}^m norm_N(r_{A_i}, s_{A_i}) \cdot r_{A_i}$$

$$\tilde{s} = s_{candidate} + \delta \cdot \alpha \sum_{i=1}^m norm_N(r_{A_i}, s_{A_i}) \cdot s_{A_i}$$

The user-specified delegation factor α defines how much base weight an experience with an associated service provider has in relation to an experience made with the candidate provider itself. δ is a scaling factor that limits the influence of delegated information as certainty increases.




Fig. 5. Coalition Forming and Verification of Cooperation Messages

$$\delta = \frac{N - (r_{candidate} + s_{candidate})}{N}$$

Specifically, under total uncertainty ($r_{candidate} + s_{candidate} = 0$) $\delta = 1$, under complete certainty ($r_{candidate} + s_{candidate} = N$) $\delta = 0$.

$c_{candidate}$ and $t_{candidate}$ are computed based on \tilde{r}, \tilde{s} instead of $r_{candidate}, s_{candidate}$.

$$c_{candidate} = \frac{N(\tilde{r} + \tilde{s})}{2(N - (\tilde{r} + \tilde{s})) + N(\tilde{r} + \tilde{s})}$$

$$t_{candidate} = \frac{\tilde{r}}{\tilde{r} + \tilde{s}}$$

Additionally, only experiences of those A_i with a certainty higher than a specific threshold might be taken into account. This would increase the impact of reputable and generally well-known coalition partners.

Thus, the expected utility for the consumer is $EU := p_{candidate} \cdot G - (1 - p_{candidate}) \cdot L$. $p_{candidate}$ is approximated as $p_{candidate} \approx E_{f,w,N}(t_{candidate}, c_{candidate}) = c_{candidate} \cdot t_{candidate} + (1 - c_{candidate}) \cdot f$.

Table 3. Reputation Updates with Coalitions

Interaction		Update	
Provider Associates		Provider Associates	
success	–	positive	see text
failure	–	negative	see text

The trust updates after an interaction can be found in table 3: only new evidence for the selected service provider is collected regarding its performance. The selected provider alone is responsible for its performance as the only influence of the associates is the association itself. The future performance of the associates is independent from the selected provider. If the service provider and the associate are not part of same service composition, new evidence for the associates is collected only in the context of their ability to reliably form association. If they are, however, part of the same composite service (cf. section 2), the reputation is updated for all service components.

4 Evaluation

4.1 Evaluation within Use Case

The use case presented in section 2 introduces a composite cyber-physical process, in which some service components/providers are visible to the users, while others are contained in a grey box internal process. We deem this use case to be typical of an online goods ordering process. The payment functionality for the service is provided through a credit card company, while the delivery is handled by an independent parcel service. The grey box process is certified by a certification provider.

Assumptions. It can reasonably be assumed that the credit card company is well-known to and trusted by the customer. This stems both from past experiences, as well as (and possibly more importantly) from strong contractual obligations between a customer and his credit card company. Similar obligations exist between the credit card company and the provider of the composite service. Thus, social and legal assurances are in place to enforce the dependability of the partners in this setting. Furthermore, because a large number of internet services use a small number of credit card companies, experience with the credit card provider generally increases more rapidly than experience with any particular composite cyber-physical service. Additionally, a credit card company within a service composition offers insurance services to its customers.

Within the use case, the grey box internal process is certified by a certification provider (ideally following a thorough and transparent audit), for instance ISO (e.g., for quality management) or TRUSTe (for privacy, however cf. [1]). We abstract from the multi-dimensionality of trust within the scope of this paper. Certification providers are less strongly coupled with a service than the aforementioned credit card company. We assume that a limited number of certification providers is used by a considerable number of services, thus easing trust establishment on certification provider. Paying for a certification by a reputable certification provider indicates a service provider's initial commitment to remaining in a market (i.e., an incentive not to defect) [3].

Both insurance and certification depend heavily on reliance [20] on a third party. Trust in the insurance and certification providers to enforce user interests in case of service provider defection has to be established. If a certification provider is incapable or unwilling to enforce its certification rigorously, a certification can actually be interpreted as a sign of untrustworthiness [1]. It is therefore assumed that the user can reliably establish trust on insurance and certification providers using a trust model.

The shipping service represents the physical interface of the composite service to the customer. While the reliability of the shipping provider is essential to a successful overall service provisioning, it is not strongly coupled to the grey box internal process of the use case.

Component Integration. Modeling overall reliability trust in the unknown service composition requires combining the information on its components. Due to the highly regulated relationship between the the credit card provider and the grey box internal component of the service composition, the providers of these two components are considered to be in a coalition (cf. section 3.5). Therefore, the well-established trust the users has in its credit card provider is delegated to the internal component. As the shipping service is essential to the success of an interaction between customer and the service composition, but is only relatively loosely coupled to it, we propose the use of the CertainLogic AND operator (\wedge_{CL}) [23]. Including a certification provider to certify the grey box internal process (for which no prior experience has been recorded), the overall computed reliability trust in the unknown composite $c_{composite} = 0$ thus becomes:

$$p_{composite} \approx (t_{credit} \cdot \alpha \cdot c_{credit} + (1 - c_{credit} \cdot (f_{cert}))) \wedge_{CL} E(t_{shipping}, c_{shipping})$$

Under a complete lack of information on *any* part of the composite service, the reliability trust value of the indicator-augmented trust computation corresponds to the CertainTrust value without indicators. The return value for $p_{composite}$ in this case is the user’s initial expectation f . The same condition holds for complete certainty, i.e., $c_{composite} = 1$, in which case $p_{composite}$ is approximated as $t_{composite}$.

Figure 6 shows the behavior of trust evaluation of CertainTrust with and without indicators over 10 interactions (for $N = 10$ and $f = 0.5$). The trustworthiness of the credit card company and the certification provider were assumed to be high ($p = 0.95$) and known to the user at this level with certainty ($c = 1$). In this way, coalition and certification was essentially used to dynamically alter the initial trust in the unknown composite service, from $f = 0.5$ for the base CertainTrust case without indicators, to 0.95. The composite cyber-physical service from our use case was therefore initially evaluated by the user at $p_{composite} \approx 0.95$. While trustworthy service providers can thereby overcome cold start issues effectively, it theoretically offers malicious service providers a considerably bigger potential to exploit this positive reputation.

The increase of the initial trust expectation from 0.5 to 0.95, however, was not arbitrary. Increasing the reliability trust in the unknown service was based on two criteria. The weaker one, certification, that the certification provider (e.g., ISO) would audit the service provider and possible revoke the certification in case of a complaint against the service. This certification provider backs this with its own behavior. The second, stronger criterion from a customer perspective, is the stronger reliance the credit card payment process p offers. Because the credit card company does not stake its reputation, but also direct monetary values through an insurance service, it has a strong incentive to actually enforce the contractual obligations between itself and the core component of the unknown service composition (the grey box).

The reliance introduced through the credit card payment process does not only justify adjusting the initial expectation value of the reliability trust upwards, but also directly influences the customer’s decision criterion, as per equation 2. This equation reflects the level of protection the credit card provider offers for an interaction with a possibly fraudulent service. For our use case, we assume that the cost of the ordered good (this includes additional costs such as shipping & handling) is paid upfront through a credit card. This money is potentially lost in the interaction, it therefore represents $L_{candidate}$. The gain G is at least as high as $L_{candidate}$, otherwise it would be unreasonable to begin the transaction. The cost of claiming a credit card insurance is assumed to be negligible compared to the cost of the

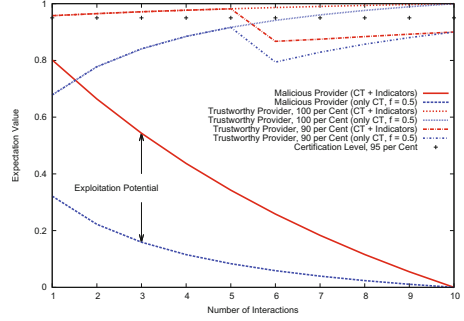


Fig. 6. Reliability Trust Expectation, for $N = 10$ and $f = 0.5$

product, while the fixed costs of the insurance ($L_{insurer}^{fix}$) are covered via a surcharge on shipping and handling levied by the service provider. Due to strong contractual agreements between the customer and the credit card company, the trustworthiness of the credit card provider (expressed as $p_{insurer}$) can be practically assured. Assuming that $L_{candidate} = G$ and $p_{insurer} \approx 1$, the decision criterion for the use case thus becomes $EU := p_{composite} \cdot G - (1 - p_{composite}) \cdot (1 - p_{insurer}) \cdot G - L_{insurer}^{fix}$. For $p_{composite} \ll 1$, as would be the case when facing an unknown service, the expected utility is considerably higher for the insurance through credit card case than it would be without the insurance option. Thus, even under the risk of increasing the exploitation potential w.r.t. malicious service providers, reliance mechanisms still allow the customer to feel safe.

4.2 Simulation

In order to show the feasibility of the proposed mechanisms in a qualitative way, each was implemented in the agent-based simulation framework used in [7]. The basic CertainTrust trust model [22] was used for evaluating providers, using $E_{f=0.5, w=1, N=10}$. The decision criterion was expected utility, as outlined in the previous chapters, with *softmax* and a decaying temperature parameter. A consumer population of 250 agents was arrayed in a clustered social network (generated according to [10]), to serve as recommenders. The same basic configuration was used to test all mechanisms against a base case, solely using experience and witness recommendation to select providers. The market was started with 15 providers (5 with $0.8 < p_{candidate} \leq 0.95$, 5 with $0.5 < p_{candidate} \leq 0.8$ and 5 with $0 < p_{candidate} \leq 0.5$) and ran for 800 rounds. At round 300, a new provider with $p_{candidate} = 0.95$ is added, in order to test the market entry performance of the different mechanisms. The objective is for the consumers to select the best provider by learning their trustworthiness.

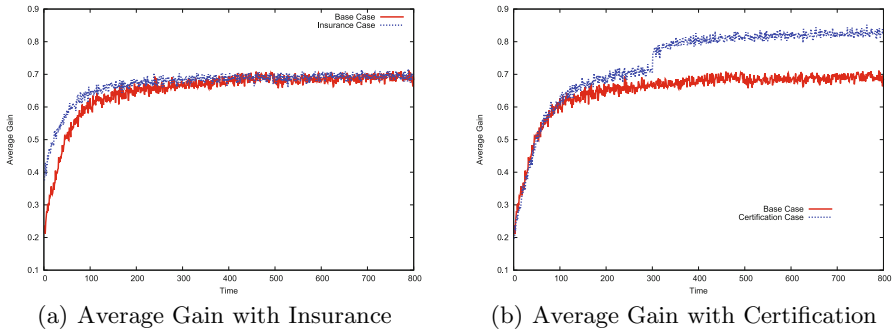


Fig. 7. Agent-based Simulation Results for Insurance and Certification

Insurance. As figure 7(a) shows, over the entire simulation run, the performance of the insurance mechanism (measured as the averaged gain over all consumers) approaches the base case. Significantly better performance, as determined by a Wilcoxon rank-sum test (95 % confidence), was attained in the

initial phase of the learning process, i.e., between timesteps 0 and 250. In this early phase, the *softmax* algorithm causes a higher exploration rate, thus leading to a higher proportion of untrustworthy providers with $p_{candidate} \leq 0.5$. Losses incurred are compensated by insurance providers, represented as randomly assigned agents with $0.5 < p_{insurer} \leq 0.95$.

Certification. The effects of certification (figure 7(b)) are complementary to the insurance case. While showing no improvement over the base case in the early rounds, it facilitates easier market entry for new providers with a high trustworthiness. The certification providers are assumed to be honest and certify conservatively ($q_{cert} = p_{candidate} - 0.1$). Certifier performance was learned using the CertainTrust trust model independently. The considerable improvement at timestep 300 is caused by the addition of the new, trustworthy provider, which is selected based on its certification, despite *softmax* already being highly exploitative.

Coalitions. Coalitions outperform the base case (figure 8) after initial exploration significantly. This is caused by trustworthy providers dissolving coalitions with less trustworthy ones, leading to highly selected coalitions of good providers. For this simulation, coalitions are formed with up to 2 other providers. Each provider in a coalition operates non-competitively from its associates, i.e., the simulation was run with three different provider populations of 15 providers each. Only one such market is plotted.

5 Related Work

Reputation and trust for eCommerce, as well as other fields, such as wireless routing, p2p networks or agent-systems, has been receiving considerable attention. An increasing number of survey articles attests to this ongoing interest, e.g., [6,11,24,25]. Typically, reputation-based trust models are driven by direct experience and witness recommendations [5,8,22]. In [12], the authors argue that comprehensive (reputation-based) trust management systems have to enable users

to assess providers reliably *and* that providers have to be given the chance to represent their trustworthiness. While the former has been the focus of much of the cited work, the latter still requires considerable efforts. Some trust models, such as FIRE [9], are modular to enable the integration of additional components, beyond experience and recommendations.

In [8], the authors address the exploration-vs-exploitation dilemma in trust-based service selection explicitly. This is, however, not done by incorporating additional information, but by analyzing temporal changes in provider behavior and adjusting random exploration accordingly.




Fig. 8. Average Gain with Coalitions ($\alpha = 0.5$) Compared to Base Case

Influences of reputation to the providers' amount of interactions have been shown in [21], exemplarily for eBay. The impact of reputation on revenue increases the attractiveness for attacks on reputation systems, leading to ongoing research in the design of robust reputation systems [14]. Incentivizing honest behavior has been directly linked to the ease with which providers can enter and leave a market [2,3,15].

Trust-based decision-making for eCommerce from a more user-centric perspective is formalized by [16]. They propose a conceptual framework to put trust, risk and their antecedents into context, lack however a computational integration.

6 Conclusions

We proposed three mechanism as indicators of trustworthiness for reputation-based trust metrics that influence the initial expectation of a customer towards a service. Each indicator has a distinct impact on the overall provider selection by consumer populations, allowing consumers to reduce their risk (insurance) and providers to represent their capabilities (certification and coalitions). By investing resources and staking reputation, service providers represent their commitment to a market, easing the service selection problem for the consumers. Future work will test the proposed and further indicators in a more comprehensive and quantitative manner, as well as investigating machine learning methods to predict trustworthy behavior based on (further) indicators. Empirical work on the positive and negative impact of certifications (e.g., [1]) is to be integrated into adapting initial expectations in the used trust model. Furthermore, specific trust-based exploration-vs-exploitation strategies will be integrated with indicators of trustworthiness.

Acknowledgments. The work presented in this paper was performed in the context of the Software-Cluster projects EMERGENT and InDiNet (www.software-cluster.org) and funded by the German Federal Ministry of Education and Research (BMBF) under grants no. "01IC10S01" and "01IC10S04". The authors assume responsibility for the content.

References

1. Edelman, B.: Adverse selection in online trust certifications. In: Proceedings of the 11th International Conference on Electronic Commerce, pp. 205–212. ACM (2009)
2. Feldman, M., Papadimitriou, C., Chuang, J., Stoica, I.: Free-riding and whitewashing in peer-to-peer systems. In: Proceedings of the ACM SIGCOMM Workshop on Practice and Theory of Incentives in Networked Systems, PINS 2004, pp. 228–236. ACM, New York (2004)
3. Friedman, E., Resnick, P.: The social cost of cheap pseudonyms. *Ann Arbor* 1001(2), 48109–1092 (1999)
4. Gambetta, D.: Can We Trust Trust? In: Gambetta, D. (ed.) *Trust: Making and Breaking Cooperative Relations*, pp. 213–237. Basil Blackwell, Oxford (1988)
5. Golbeck, J.: *Computing and applying trust in web-based social networks*. Doctoral thesis, University of Maryland (2005)
6. Grandison, T., Sloman, M.: A Survey of Trust in Internet Applications. *IEEE Communications and Survey* 3(4), 2–16 (2000)

7. Hauke, S., Pyka, M., Heider, D.: Towards Improved Trust Diffusion Through Active Recommender Propagation. In: 4th International Conference on Complex Distributed Systems (2010)
8. Hoogendoorn, M., Jaffry, S.W., Treur, J.: Incorporating Interdependency of Trust Values in Existing Trust Models for Trust Dynamics. In: Nishigaki, M., Jøsang, A., Murayama, Y., Marsh, S. (eds.) IFIPTM 2010. IFIP AICT, vol. 321, pp. 263–276. Springer, Heidelberg (2010)
9. Huynh, T.D.: Trust and Reputation in Open Multi-Agent Systems. PhD thesis, University of Southampton (2006)
10. Jin, E.M., Girvan, M., Newman, M.E.J.: Structure of growing social networks. *Phys. Rev. E* 64(4), 46132 (2001)
11. Jøsang, A., Ismail, R., Boyd, C.: A survey of Trust and Reputation Systems for Online Service Provision. *Decision Support Systems* 43(2), 618–644 (2007)
12. Jøsang, A., Keser, C., Dimitrakos, T.: Can We Manage Trust? In: Herrmann, P., Issarny, V., Shiu, S.C.K. (eds.) iTrust 2005. LNCS, vol. 3477, pp. 93–107. Springer, Heidelberg (2005)
13. Jøsang, A., Lo Presti, S.: Analysing the Relationship between Risk and Trust. In: Jensen, C., Poslad, S., Dimitrakos, T. (eds.) iTrust 2004. LNCS, vol. 2995, pp. 135–145. Springer, Heidelberg (2004)
14. Kerr, R., Cohen, R.: An Experimental Testbed for Evaluation of Trust and Reputation Systems. In: Ferrari, E., Li, N., Bertino, E., Karabulut, Y. (eds.) IFIPTM 2009. IFIP AICT, vol. 300, pp. 252–266. Springer, Heidelberg (2009)
15. Kerr, R., Cohen, R.: Trust as a Tradable Commodity: A Foundation for Safe Electronic Marketplaces. *Computational Intelligence* 26(2), 160–182 (2010)
16. Kim, D., Ferrin, D., Rao, H.: A trust-based consumer decision-making model in electronic commerce: The role of trust, perceived risk, and their antecedents. *Decision Support Systems* 44(2), 544–564 (2008)
17. Luce, R.D.: Utility of gains and losses: Measurement-theoretical and experimental approaches. Lawrence Erlbaum Associates Publishers, Mahwah (2000)
18. Marsh, S., Dibben, M.R.: Trust, Untrust, Distrust and Mistrust – An Exploration of the Dark(er) Side. In: Herrmann, P., Issarny, V., Shiu, S. (eds.) iTrust 2005. LNCS, vol. 3477, pp. 17–33. Springer, Heidelberg (2005)
19. McKnight, D.H., Chervany, N.L.: Trust and Distrust Definitions: One Bite at a Time. In: Falcone, R., Singh, M., Tan, Y.-H. (eds.) Trust in Cyber-societies. LNCS (LNAI), vol. 2246, pp. 27–54. Springer, Heidelberg (2001)
20. Pichler, R.: Trust and Reliance - Enforcement and Compliance: Enhancing Consumer Confidence in the Electronic Marketplace. Juridical sciences master, Stanford University (2000)
21. Resnick, P., Zeckhauser, R., Swanson, J., Lockwood, K.: The value of reputation on ebay: A controlled experiment. *Experimental Economics* 9, 79–101 (2003)
22. Ries, S.: Trust in Ubiquitous Computing. Doctoral thesis, TU Darmstadt (2009)
23. Ries, S., Habib, S.M., Mühlhäuser, M., Varadharajan, V.: CertainLogic: A Logic for Modeling Trust and Uncertainty (short paper). In: McCune, J.M., Balacheff, B., Perrig, A., Sadeghi, A.-R., Sasse, A., Beres, Y. (eds.) Trust 2011. LNCS, vol. 6740, pp. 254–261. Springer, Heidelberg (2011)
24. Sabater, J., Sierra, C.: Review on Computational Trust and Reputation Models. *Artificial Intelligence Review* 24, 33–60 (2005)
25. Wang, Y., Vassileva, J.: Toward trust and reputation based web service selection: A survey. *International Transactions on Systems Science and Applications* 3(2), 118–132 (2007)