Rumours and Reputation: Evaluating Multi-Dimensional Trust within a Decentralised Reputation System

Steven Reece[1], Alex Rogers[2], Stephen Roberts[1] and Nicholas R. Jennings[2]

[1] Department of Engineering Science, University of Oxford, Oxford, OX1 3PJ, UK.
{reece,sjrob}@robots.ox.ac.uk

[2] Electronics and Computer Science, University of Southampton, Southampton, SO17 1BJ, UK.
{acr,nrj}@ecs.soton.ac.uk

ABSTRACT

In this paper we develop a novel probabilistic model of computational trust that explicitly deals with correlated multi-dimensional contracts. Our starting point is to consider an agent attempting to estimate the utility of a contract, and we show that this leads to a model of computational trust whereby an agent must determine a vector of estimates that represent the probability that any dimension of the contract will be successfully fulfilled, and a covariance matrix that describes the uncertainty and correlations in these probabilities. We present a formalism based on the Dirichlet distribution that allows an agent to calculate these probabilities and correlations from their direct experience of contract outcomes, and we show that this leads to superior estimates compared to an alternative approach using multiple independent beta distributions. We then show how agents may use the sufficient statistics of this Dirichlet distribution to communicate and fuse reputation within a decentralised reputation system. Finally, we present a novel solution to the problem of rumour propagation within such systems. This solution uses the notion of private and shared information, and provides estimates consistent with a centralised reputation system, whilst maintaining the anonymity of the agents, and avoiding bias and overconfidence.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Intelligent agents

General Terms

Algorithms, Design, Theory

Keywords

multi-dimensional trust, rumour propogation, Dirichlet distribution

1. INTRODUCTION

The role of computational models of trust within multi-agent systems in particular, and open distributed systems in general, has recently generated a great deal of research interest. In such systems, agents must typically choose between interaction partners, and in this context trust can be viewed to provide a means for agents to represent and estimate the reliability with which these interaction partners will fulfill their commitments. To date, however, much of the work within this area has used domain specific or ad hoc trust metrics, and has focused on providing heuristics to evaluate and update these metrics using direct experience and reputation reports from other agents (see [8] for a review).

Recent work has attempted to place the notion of computational trust within the framework of probability theory [6, 11]. This approach allows many of the desiderata of computational trust models to be addressed through principled means. In particular: (i) it allows agents to update their estimates of the trustworthiness of a supplier as they acquire direct experience, (ii) it provides a natural framework for agents to express their uncertainty this trustworthiness, and, (iii) it allows agents to exchange, combine and filter reputation reports received from other agents.

Whilst this approach is attractive, it is somewhat limited in that it has so far only considered single dimensional outcomes (i.e. whether the contract has succeeded or failed in its entirety). However, in many real world settings the success or failure of an interaction may be decomposed into several dimensions [7]. This presents the challenge of combining these multiple dimensions into a single metric over which a decision can be made. Furthermore, these dimensions will typically also exhibit correlations. For example, a contract within a supply chain may specify criteria for timeliness, quality and quantity. A supplier who is suffering delays may attempt a trade-off between these dimensions by supplying the full amount late, or supplying as much as possible (but less than the quantity specified within the contract) on time. Thus, correlations will naturally arise between these dimensions, and hence, between the probabilities that describe the successful fulfillment of each contract dimension. To date, however, no such principled framework exists to describe these multi-dimensional contracts, nor the correlations between these dimensions (although some ad hoc models do exist – see section 2 for more details).

To rectify this shortcoming, in this paper we develop a probabilistic model of computational trust that explicitly deals with correlated multi-dimensional contracts. The starting point for our work is to consider how an agent can estimate the utility that it will derive from interacting with a supplier. Here we use standard approaches from the literature of data fusion (since this is a well developed field where the notion of multi-dimensional correlated estimates is well established[1]) to show that this naturally leads to a trust model where the agent must estimate probabilities and correlations over

[1] In this context, the multiple dimensions typically represent the physical coordinates of a target being tracked, and correlations arise through the operation and orientation of sensors.

AAMAS'07 May 14–18 2007, Honolulu, Hawai'i, USA.

multiple dimensions. Building upon this, we then devise a novel trust model that addresses the three desiderata discussed above. In more detail, in this paper we extend the state of the art in four key ways:

1. We devise a novel multi-dimensional probabilistic trust model that enables an agent to estimate the expected utility of a contract, by estimating (i) the probability that each contract dimension will be successfully fulfilled, and (ii) the correlations between these estimates.
2. We present an exact probabilistic model based upon the Dirichlet distribution that allows agents to use their direct experience of contract outcomes to calculate the probabilities and correlations described above. We then benchmark this solution and show that it leads to good estimates.
3. We show that agents can use the sufficient statistics of this Dirichlet distribution in order to exchange reputation reports with one another. The sufficient statistics represent aggregations of their direct experience, and thus, express contract outcomes in a compact format with no loss of information.
4. We show that, while being efficient, the aggregation of contract outcomes can lead to double counting, and rumour propagation, in decentralised reputation systems. Thus, we present a novel solution based upon the idea of private and shared information. We show that it yields estimates consistent with a centralised reputation system, whilst maintaining the anonymity of the agents, and avoiding overconfidence.

The remainder of this paper is organised as follows: in section 2 we review related work. In section 3 we present our notation for a single dimensional contract, before introducing our multi-dimensional trust model in section 4. In sections 5 and 6 we discuss communicating reputation, and present our solution to rumour propagation in decentralised reputation systems. We conclude in section 7.

2. RELATED WORK

The need for a multi-dimensional trust model has been recognised by a number of researchers. Sabater and Sierra present a model of reputation, in which agents form contracts based on multiple variables (such as delivery date and quality), and define *impressions* as subjective evaluations of the outcome of these contracts. They provide heuristic approaches to combining these impressions to form a measure they call subjective reputation.

Likewise, Griffiths decomposes overall trust into a number of different dimensions such as success, cost, timeliness and quality [4]. In his case, each dimension is scored as a real number that represents a comparative value with no strong semantic meaning. He develops an heuristic rule to update these values based on the direct experiences of the individual agent, and an heuristic function that takes the individual trust dimensions and generates a single scalar that is then used to select between suppliers. Whilst, he comments that the trust values could have some associated confidence level, heuristics for updating these levels are not presented.

Gujral et al. take a similar approach and present a trust model over multiple domain specific dimensions [5]. They define multi-dimensional goal requirements, and evaluate an expected payoff based on a supplier's estimated behaviour. These estimates are, however, simple aggregations over the direct experience of several agents, and there is no measure of the uncertainty. Nevertheless, they show that agents who select suppliers based on these multiple dimensions outperform those who consider just a single one.

By contrast, a number of researchers have presented more principled computational trust models based on probability theory, albeit limited to a single dimension. Jøsang and Ismail describe the Beta Reputation System whereby the reputation of an agent is compiled from the positive and negative reports from other agents who have interacted with it [6]. The beta distribution represents a natural choice for representing these binary outcomes, and it provides a principled means of representing uncertainty. Moreover, they provide a number of extensions to this initial model including an approach to exchanging reputation reports using the sufficient statistics of the beta distribution, methods to discount the opinions of agents who themselves have low reputation ratings, and techniques to deal with reputations that may change over time.

Likewise, Teacy et al. use the beta distribution to describe an agent's belief in the probability that another agent will successfully fulfill its commitments [11]. They present a formalism using a beta distribution that allows the agent to estimate this probability based upon its direct experience, and again they use the sufficient statistics of this distribution to communicate this estimate to other agents. They provide a number of extensions to this initial model, and, in particular, they consider that agents may not always truthfully report their trust estimates. Thus, they present a principled approach to detecting and removing inconsistent reports.

Our work builds upon these more principled approaches. However, the starting point of our approach is to consider an agent that is attempting to estimate the expected utility of a contract. We show that estimating this expected utility requires that an agent must estimate the probability with which the supplier will fulfill its contract. In the single-dimensional case, this naturally leads to a trust model using the beta distribution (as per Jøsang and Ismail and Teacy et al.). However, we then go on to extend this analysis to multiple dimensions, where we use the natural extension of the beta distribution, namely the Dirichlet distribution, to represent the agent's belief over multiple dimensions.

3. SINGLE-DIMENSIONAL TRUST

Before presenting our multi-dimensional trust model, we first introduce the notation and formalism that we will use by describing the more familiar single dimensional case. We consider an agent who must decide whether to engage in a future contract with a supplier. This contract will lead to some outcome, o, and we consider that $o = 1$ if the contract is successfully fulfilled, and $o = 0$ if not[2].

In order for the agent to make a rational decision, it should consider the utility that it will derive from this contract. We assume that in the case that the contract is successfully fulfilled, the agent derives a utility $u(o = 1)$, otherwise it receives no utility[3]. Now, given that the agent is uncertain of the reliability with which the supplier will fulfill the contract, it should consider the expected utility that it will derive, $E[U]$, and this is given by:

$$E[U] = p(o = 1)u(o = 1) \tag{1}$$

where $p(o = 1)$ is the probability that the supplier will successfully fulfill the contract. However, whilst $u(o = 1)$ is known by the agent, $p(o = 1)$ is not. The best the agent can do is to determine a distribution over possible values of $p(o = 1)$ given its direct experience of previous contract outcomes. Given that it has been able to do so, it can then determine an estimate of the expected utility[4] of the contract, $E[E[U]]$, and a measure of its uncertainty in this expected utility, Var($E[U]$). This uncertainty is important since a risk averse agent may make a decision regarding a contract,

[2]Note that we only consider binary contract outcomes, although extending this to partial outcomes is part of our future work.

[3]Clearly this can be extended to the case where some utility is derived from an unsuccessful outcome.

[4]Note that this is often called the "expected expected utility", and this is the notation that we adopt here [2].

not only on its estimate of the expected utility of the contract, but also on the probability that the expected utility will exceed some minimum amount. These two properties are given by:

$$E[E[U]] = \hat{p}(o = 1)u(o = 1) \tag{2}$$

$$\text{Var}(E[U]) = \text{Var}(p(o = 1))u(o = 1)^2 \tag{3}$$

where $\hat{p}(o = 1)$ and $\text{Var}(p(o = 1))$ are the estimate and uncertainty of the probability that a contract will be successfully fulfilled, and are calculated from the distribution over possible values of $p(o = 1)$ that the agent determines from its direct experience. The utility based approach that we present here provides an attractive motivation for this model of Teacy et al. [11].

Now, in the case of binary contract outcomes, the beta distribution is the natural choice to represent the distribution over possible values of $p(o = 1)$ since within Bayesian statistics this well known to be the conjugate prior for binomial observations [3]. By adopting the beta distribution, we can calculate $\hat{p}(o = 1)$ and $\text{Var}(p(o = 1))$ using standard results, and thus, if an agent observed N previous contracts of which n were successfully fulfilled, then:

$$\hat{p}(o = 1) = \frac{n + 1}{N + 2}$$

and:

$$\text{Var}(p(o = 1)) = \frac{(n + 1)(N - n + 1)}{(N + 2)^2(N + 3)}$$

Note that as expected, the greater the number of contracts the agent observes, the smaller the variance term $\text{Var}(p(o = 1))$, and, thus, the less the uncertainty regarding the probability that a contract will be successfully fulfilled, $\hat{p}(o = 1)$.

4. MULTI-DIMENSIONAL TRUST

We now extend the description above, to consider contracts between suppliers and agents that are represented by multiple dimensions, and hence the success or failure of a contract can be decomposed into the success or failure of each separate dimension. Consider again the example of the supply chain that specifies the timeliness, quantity, and quality of the goods that are to be delivered. Thus, within our trust model "$o_a = 1$" now indicates a successful outcome over dimension a of the contract and "$o_a = 0$" indicates an unsuccessful one. A contract outcome, X, is now composed of a vector of individual contract part outcomes (e.g. $X = \{o_a = 1, o_b = 0, o_c = 0, \ldots\}$).

Given a multi-dimensional contract whose outcome is described by the vector X, we again consider that in order for an agent to make a rational decision, it should consider the utility that it will derive from this contract. To this end, we can make the general statement that the expected utility of a contract is given by:

$$E[U] = p(X)U(X)^T \tag{4}$$

where $p(X)$ is a joint probability distribution over all possible contract outcomes:

$$p(X) = \begin{pmatrix} p(o_a = 1, o_b = 0, o_c = 0, \ldots) \\ p(o_a = 1, o_b = 1, o_c = 0, \ldots) \\ p(o_a = 0, o_b = 1, o_c = 0, \ldots) \\ \vdots \end{pmatrix} \tag{5}$$

and $U(X)$ is the utility derived from these possible outcomes:

$$U(X) = \begin{pmatrix} u(o_a = 1, o_b = 0, o_c = 0, \ldots) \\ u(o_a = 1, o_b = 1, o_c = 0, \ldots) \\ u(o_a = 0, o_b = 1, o_c = 0, \ldots) \\ \vdots \end{pmatrix} \tag{6}$$

As before, whilst $U(X)$ is known to the agent, the probability distribution $p(X)$ is not. Rather, given the agent's direct experience of the supplier, the agent can determine a distribution over possible values for $p(X)$. In the single dimensional case, a beta distribution was the natural choice over possible values of $p(o = 1)$. In the multi-dimensional case, where $p(X)$ itself is a vector of probabilities, the corresponding natural choice is the Dirichlet distribution, since this is a conjugate prior for multinomial proportions [3].

Given this distribution, the agent is then able to calculate an estimate of the expected utility of a contract. As before, this estimate is itself represented by an expected value given by:

$$E[E[U]] = \hat{p}(X)U(X)^T \tag{7}$$

and a variance, describing the uncertainty in this expected utility:

$$\text{Var}(E[U]) = U(X)\text{Cov}(p(X))U(X)^T \tag{8}$$

where:

$$\text{Cov}(p(X)) \triangleq E[(p(X) - \hat{p}(X))(p(X) - \hat{p}(X))^T] \tag{9}$$

Thus, whilst the single dimensional case naturally leads to a trust model in which the agents attempt to derive an estimate of probability that a contract will be successfully fulfilled, $\hat{p}(o = 1)$, along with a scalar variance that describes the uncertainty in this probability, $\text{Var}(p(o = 1))$, in this case, the agents must derive an estimate of a vector of probabilities, $\hat{p}(X)$, along with a covariance matrix, $\text{Cov}(p(X))$, that represents the uncertainty in $p(X)$ given the observed contractual outcomes. At this point, it is interesting to note that the estimate in the single dimensional case, $\hat{p}(o = 1)$, has a clear semantic meaning in relation to trust; it is the agent's belief in the probability of a supplier successfully fulfilling a contract. However, in the multi-dimensional case the agent must determine $\hat{p}(X)$, and since this describes the probability of all possible contract outcomes, including those that are completely un-fulfilled, this direct semantic interpretation is not present. In the next section, we describe the exemplar utility function that we shall use in the remainder of this paper.

4.1 Exemplar Utility Function

The approach described so far is completely general, in that it applies to any utility function of the form described above, and also applies to the estimation of any joint probability distribution. In the remainder of this paper, for illustrative purposes, we shall limit the discussion to the simplest possible utility function that exhibits a dependence upon the correlations between the contract dimensions. That is, we consider the case that expected utility is dependent only on the marginal probabilities of each contract dimension being successfully fulfilled, rather than the full joint probabilities:

$$U(X) = \begin{pmatrix} u(o_a = 1) \\ u(o_b = 1) \\ u(o_c = 1) \\ \vdots \end{pmatrix} \tag{10}$$

Thus, $\hat{p}(X)$ is a vector estimate of the probability of each contract dimension being successfully fulfilled, and maintains the clear semantic interpretation seen in the single dimensional case:

$$\hat{p}(X) = \begin{pmatrix} \hat{p}(o_a = 1) \\ \hat{p}(o_b = 1) \\ \hat{p}(o_c = 1) \\ \vdots \end{pmatrix} \tag{11}$$

The correlations between the contract dimensions affect the uncertainty in the expected utility. To see this, consider the covariance matrix that describes this uncertainty, Cov(p(X)), is now given by:

$$\text{Cov}(p(X)) = \begin{pmatrix} V_a & C_{ab} & C_{ac} & \dots \\ C_{ab} & V_b & C_{bc} & \dots \\ C_{ac} & C_{bc} & V_c & \dots \\ \vdots & \vdots & \vdots & \end{pmatrix} \tag{12}$$

In this matrix, the "diagonal" terms, V_a, V_b and V_c, represent the uncertainties in $p(o_a = 1)$, $p(o_b = 1)$ and $p(o_c = 1)$ within $p(X)$. The "off-diagonal" terms, C_{ab}, C_{ac} and C_{bc}, represent the correlations between these probabilities. In the next section, we use the Dirichlet distribution to calculate both $\hat{p}(X)$ and Cov$(p(X))$ from an agent's direct experience of previous contract outcomes.

We first illustrate why this is necessary by considering an alternative approach to modelling multi-dimensional contracts whereby an agent naïvely assumes that the dimensions are independent, and thus, it models each individually by separate beta distributions (as in the single dimensional case we presented in section 3). This is actually equivalent to setting the off-diagonal terms within the covariance matrix, Cov$(p(X))$, to zero. However, doing so can lead an agent to assume that its estimate of the expected utility of the contract is more accurate than it actually is. To illustrate this, consider a specific scenario with the following values: $u(o_a = 1) = u(o_b = 1) = 1$ and $V_a = V_b = 0.2$. In this case, Var$(E[U]) = 0.4(1 + C_{ab})$, and thus, if the correlation C_{ab} is ignored then the variance in the expected utility is 0.4. However, if the contract outcomes are completely correlated then $C_{ab} = 1$ and Var$(E[U])$ is actually 0.8. Thus, in order to have an accurate estimate of the variance of the expected contract utility, and to make a rational decision, it is essential that the agent is able to represent and calculate these correlation terms. In the next section, we describe how an agent may do so using the Dirichlet distribution.

4.2 The Dirichlet Distribution

In this section, we describe how the agent may use its direct experience of previous contracts, and the standard results of the Dirichlet distribution, to determine an estimate of the probability that each contract dimension will be successful fulfilled, $\hat{p}(X)$, and a measure of the uncertainties in these probabilities that expresses the correlations between the contract dimensions, Cov$(p(X))$.

We first consider the calculation of $\hat{p}(X)$ and the diagonal terms of the covariance matrix Cov$(p(X))$. As described above, the derivation of these results is identical to the case of the single dimensional beta distribution, where out of N contract outcomes, n are successfully fulfilled. In the multi-dimensional case, however, we have a vector $\{n_a, n_b, n_c, \ldots\}$ that represents the number of outcomes for which each of the individual contract dimensions were successfully fulfilled. Thus, in terms of the standard Dirichlet parameters where $\alpha_a = n_a + 1$ and $\alpha_0 = N + 2$, the agent can estimate the probability of this contract dimension being successfully fulfilled:

$$\hat{p}(o_a = 1) = \frac{\alpha_a}{\alpha_0} = \frac{n_a + 1}{N + 2}$$

and can also calculate the variance in any contract dimension:

$$V_a = \frac{\alpha_a(\alpha_0 - \alpha_a)}{\alpha_0^2(1 + \alpha_0)} = \frac{(n_a + 1)(N - n_a + 1)}{(N + 2)^2(N + 3)}$$

However, calculating the off-diagonal terms within Cov$(p(X))$ is more complex since it is necessary to consider the correlations between the contract dimensions. Thus, for each pair of dimensions (i.e. a and b), we must consider all possible combinations of contract outcomes, and thus we define n_{ij}^{ab} as the number of contract outcomes for which both $o_a = i$ and $o_b = j$. For example, n_{10}^{ab} represents the number of contracts for which $o_a = 1$ and $o_b = 0$.

Now, using the standard Dirichlet notation, we can define $\alpha_{ij}^{ab} \triangleq n_{ij}^{ab} + 1$ for all i and j taking values 0 and 1, and then, to calculate the cross-correlations between contract pairs a and b, we note that the Dirichlet distribution over pair-wise joint probabilities is:

$$\text{Prob}(p_{ab}) = K_{ab} \prod_{i\in\{0,1\}} \prod_{j\in\{0,1\}} p(o_a = i, o_b = j)^{\alpha_{ij}^{ab} - 1}$$

where:

$$\sum_{i\in\{0,1\}} \sum_{j\in\{0,1\}} p(o_a = i, o_b = j) = 1$$

and K_{ab} is a normalising constant [3]. From this we can derive pair-wise probability estimates and variances:

$$E[p(o_a = i, o_b = j)] = \frac{\alpha_{ij}^{ab}}{\alpha_0} \tag{13}$$

$$V[p(o_a = i, o_b = j)] = \frac{\alpha_{ij}^{ab}(\alpha_0 - \alpha_{ij}^{ab})}{\alpha_0^2(1 + \alpha_0)} \tag{14}$$

where:

$$\alpha_0 = \sum_{i\in\{0,1\}} \sum_{j\in\{0,1\}} \alpha_{ij}^{ab} \tag{15}$$

and in fact, $\alpha_0 = N + 2$, where N is the total number of contracts observed. Likewise, we can express the covariance in these pair-wise probabilities in similar terms:

$$\text{C}[p(o_a = i, o_b = j), p(o_a = m, o_b = n)] = \frac{-\alpha_{ij}^{ab}\alpha_{mn}^{ab}}{\alpha_0^2(1 + \alpha_0)}$$

Finally, we can use the expression:

$$p(o_a = 1) = \sum_{j\in\{0,1\}} p(o_a = 1, o_b = j)$$

to determine the covariance C_{ab}. To do so, we first simplify the notation by defining $V_{ij}^{ab} \triangleq V[p(o_a = i, o_b = j)]$ and $C_{ijmn}^{ab} \triangleq C[p(o_a = i, o_b = j), p(o_a = m, o_b = n)]$. The covariance for the probability of positive contract outcomes is then the covariance between $\sum_{j\in\{0,1\}} p(o_a = 1, o_b = j)$ and $\sum_{i\in\{0,1\}} p(o_a = i, o_b = 1)$, and thus:

$$C_{ab} = C_{1001}^{ab} + C_{1101}^{ab} + C_{1011}^{ab} + V_{11}^{ab} .$$

Thus, given a set of contract outcomes that represent the agent's previous interactions with a supplier, we may use the Dirichlet distribution to calculate the mean and variance of the probability of any contract dimension being successfully fulfilled (i.e. $\hat{p}(o_a = 1)$ and V_a). In addition, by a somewhat more complex procedure we can also calculate the correlations between these probabilities (i.e. C_{ab}). This allows us to calculate an estimate of the probability that any contract dimension will be successfully fulfilled, $\hat{p}(X)$, and also represent the uncertainty and correlations in these probabilities by the covariance matrix, Cov$(p(X))$. In turn, these results may be used to calculate the estimate and uncertainty in the expected utility of the contract. In the next section we present empirical results that show that in practise this formalism yields significant improvements in these estimates compared to the naïve approximation using multiple independent beta distributions.

4.3 Empirical Comparison

In order to evaluate the effectiveness of our formalism, and show the importance of the off-diagonal terms in Cov$(p(X))$, we compare two approaches:

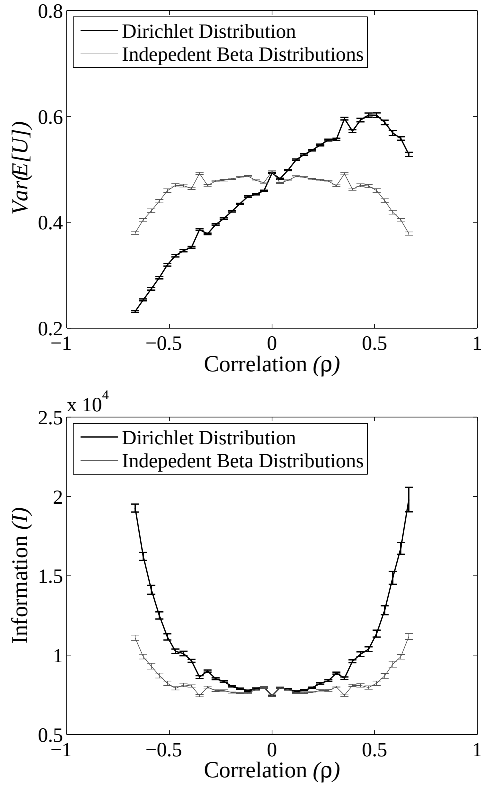

Figure 1: Plots showing (i) the variance of the expected contract utility and (ii) the information content of the estimates computed using the Dirichlet distribution and multiple independent beta distributions. Results are averaged over 10^6 runs, and the error bars show the standard error in the mean.

- **Dirichlet Distribution:** We use the full Dirichlet distribution, as described above, to calculate $\hat{p}(X)$ and Cov$(p(X))$ including all its off-diagonal terms that represent the correlations between the contract dimensions.
- **Independent Beta Distributions:** We use independent beta distributions to represent each contract dimension, in order to calculate $\hat{p}(X)$, and then, as described earlier, we approximate Cov$(p(X))$ and ignore the correlations by setting all the off-diagonal terms to zero.

We consider a two-dimensional case where $u(o_a = 1) = 6$ and $u(o_b = 1) = 2$, since this allows us to plot $\hat{p}(X)$ and Cov$(p(X))$ as ellipses in a two-dimensional plane, and thus explain the differences between the two approaches. Specifically, we initially allocate the agent some previous contract outcomes that represents its direct experience with a supplier. The number of contracts is drawn uniformly between 10 and 20, and the actual contract outcomes are drawn from an arbitrary joint distribution intended to induce correlations between the contract dimensions. For each set of contracts, we use the approaches described above to calculate $\hat{p}(X)$ and Cov$(p(X))$, and hence, the variance in the expected contract utility, Var$(E[U])$. In addition, we calculate a scalar measure of the information content, I, of the covariance matrix Cov$(p(X))$, which is a standard way of measuring the uncertainty encoded within the covariance matrix [1]. More specifically, we calculate the determinant of the inverse of the covariance matrix:

$$I = \det(\text{Cov}(p(X))^{-1}) \quad (16)$$

and note that the larger the information content, the more precise $\hat{p}(X)$ will be, and thus, the better the estimate of the expected utility that the agent is able to calculate. Finally, we use the results presented in section 4.2 to calculate the actual correlation, ρ, associated with this particular set of contract outcomes:

$$\rho = \frac{C_{ab}}{\sqrt{V_a V_b}} \quad (17)$$

where C_{ab}, V_a and V_b are calculated as described in section 4.2.

The results of this analysis are shown in figure 1. Here we show the values of I and Var$(E[U])$ calculated by the agents, plotted against the correlation of the contract outcomes, ρ, that constituted their direct experience. The results are averaged over 10^6 simulation runs. Note that as expected, when the dimensions of the contract outcomes are uncorrelated (i.e. $\rho = 0$), then both approaches give the same results. However, the value of using our formalism with the full Dirichlet distribution is shown when the correlation between the dimensions increases (either negatively or positively). As can be seen, if we approximate the Dirichlet distribution with multiple independent beta distributions, all of the correlation information contained within the covariance matrix, Cov$(p(X))$, is lost, and thus, the information content of the matrix is much lower. The loss of this correlation information leads the variance of the expected utility of the contract to be incorrect (either over or under estimated depending on the correlation)[5], with the exact amount of mis-estimation depending on the actual utility function chosen (i.e. the values of $u(o_a = 1)$ and $u(o_b = 1)$).

In addition, in figure 2 we illustrate an example of the estimates calculated through both methods, for a single exemplar set of contract outcomes. We represent the probability estimates, $\hat{p}(X)$, and the covariance matrix, Cov$(p(X))$, in the standard way as an ellipse [1]. That is, $\hat{p}(X)$ determines the position of the center of the ellipse, Cov$(p(X))$ defines its size and shape. Note that whilst the ellipse resulting from the full Dirichlet formalism accurately reflects the true distribution (samples of which are plotted as points), that calculated by using multiple independent Beta distributions (and thus ignoring the correlations) results in a much larger ellipse that does not reflect the true distribution. The larger size of this ellipse is a result of the off-diagonal terms of the covariance matrix being set to zero, and corresponds to the agent miscalculating the uncertainty in the probability of each contract dimension being fulfilled. This, in turn, leads it to miscalculate the uncertainty in the expected utility of a contract (shown in figure 1 as Var$(E[U])$).

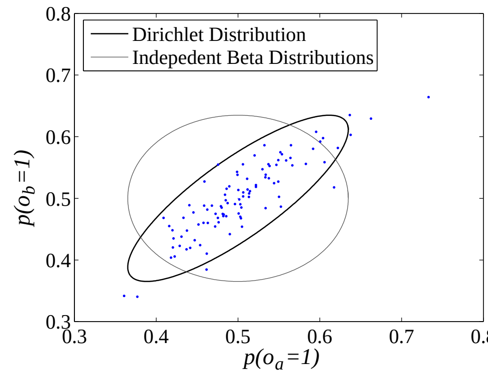

Figure 2: Examples of $\hat{p}(X)$ and Cov$(p(X))$ plotted as second standard error ellipses.

5. COMMUNICATING REPUTATION

Having described how an individual agent can use its own direct experience of contract outcomes in order to estimate the probabil-

[5] Note that the plots are not smooth due to the fact that given a limited number of contract outcomes, then the mean of V_a and V_b do not vary smoothly with ρ.

ity that a multi-dimensional contract will be successfully fulfilled, we now go on to consider how agents within an open multi-agent system can communicate these estimates to one another. This is commonly referred to as *reputation* and allows agents with limited direct experience of a supplier to make rational decisions.

Both Jøsang and Ismail, and Teacy et al. present models whereby reputation is communicated between agents using the sufficient statistics of the beta distribution [6, 11]. This approach is attractive since these sufficient statistics are simple aggregations of contract outcomes (more precisely, they are simply the total number of contracts observed, N, and the number of these that were successfully fulfilled, n). Under the probabilistic framework of the beta distribution, reputation reports in this form may simply be aggregated with an agent's own direct experience, in order to gain a more precise estimate based on a larger set of contract outcomes.

We can immediately extend this approach to the multi-dimensional case considered here, by requiring that the agents exchange the sufficient statistics of the Dirichlet distribution instead of the beta distribution. In this case, for each pair of dimensions (i.e. a and b), the agents must communicate a vector of contract outcomes, $\mathcal{N}$, which are the sufficient statistics of the Dirichlet distribution, given by:

$$\mathcal{N} = < n_{ij}^{ab} > \quad \forall a, b, i \in \{0, 1\}, j \in \{0, 1\} \tag{18}$$

Thus, an agent is able to communicate the sufficient statistics of its own Dirichlet distribution in terms of just $2d(d-1)$ numbers (where d is the number of contract dimensions). For instance, in the case of three dimensions, $\mathcal{N}$, is given by:

$$\mathcal{N} = < n_{00}^{ab}, n_{01}^{ab}, n_{10}^{ab}, n_{11}^{ab}, n_{00}^{ac}, n_{01}^{ac}, n_{10}^{ac}, n_{11}^{ac}, n_{00}^{bc}, n_{01}^{bc}, n_{10}^{bc}, n_{11}^{bc} >$$

and, hence, large sets of contract outcomes may be communicated within a relatively small message size, with no loss of information. Again, agents receiving these sufficient statistics may simply aggregate them with their own direct experience in order to gain a more precise estimate of the trustworthiness of a supplier.

Finally, we note that whilst it is not the focus of our work here, by adopting the same principled approach as Jøsang and Ismail, and Teacy et al., many of the techniques that they have developed (such as discounting reports from unreliable agents, and filtering inconsistent reports from selfish agents) may be directly applied within this multi-dimensional model. However, we now go on to consider a new issue that arises in both the single and multi-dimensional models, namely the problems that arise when such aggregated sufficient statistics are propagated within decentralised agent networks.

6. RUMOUR PROPAGATION WITHIN REPUTATION SYSTEMS

In the previous section, we described the use of sufficient statistics to communicate reputation, and we showed that by aggregating contract outcomes together into these sufficient statistics, a large number of contract outcomes can be represented and communicated in a compact form. Whilst, this is an attractive property, it can be problematic in practise, since the individual provenance of each contract outcome is lost in the aggregation. Thus, to ensure an accurate estimate, the reputation system must ensure that each observation of a contract outcome is included within the aggregated statistics no more than once.

Within a centralised reputation system, where all agents report their direct experience to a trusted center, such double counting of contract outcomes is easy to avoid. However, in a decentralised reputation system, where agents communicate reputation to one another, and aggregate their direct experience with these reputation reports on-the-fly, avoiding double counting is much more difficult.

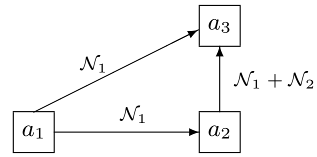

Figure 3: Example of rumour propagation in a decentralised reputation system.

For example, consider the case shown in figure 3 where three agents ($a_1 \ldots a_3$), each with some direct experience of a supplier, share reputation reports regarding this supplier. If agent a_1 were to provide its estimate to agents a_2 and a_3 in the form of the sufficient statistics of its Dirichlet distribution, then these agents can aggregate these contract outcomes with their own, and thus obtain more precise estimates. If at a later stage, agent a_2 were to send its aggregate vector of contract outcomes to agent a_3, then agent a_3 being unaware of the full history of exchanges, may attempt to combine these contract outcomes with its own aggregated vector. However, since both vectors contain a contribution from agent a_1, these will be counted twice in the final aggregated vector, and will result in a biased and overconfident estimate. This is termed *rumour propagation* or *data incest* in the data fusion literature [9].

One possible solution would be to uniquely identify the source of each contract outcome, and then propagate each vector, along with its label, through the network. Agents can thus identify identical observations that have arrived through different routes, and after removing the duplicates, can aggregate these together to form their estimates. Whilst this appears to be attractive in principle, for a number of reasons, it is not always a viable solution in practise [12]. Firstly, agents may not actually wish to have their uniquely labelled contract outcomes passed around an open system, since such information may have commercial or practical significance that could be used to their disadvantage. As such, agents may only be willing to exchange identifiable contract outcomes with a small number of other agents (perhaps those that they have some sort of reciprocal relationship with). Secondly, the fact that there is no aggregation of the contract outcomes as they pass around the network means that the message size increases over time, and the ultimate size of these messages is bounded only by the number of agents within the system (possibly an extremely large number for a global system). Finally, it may actually be difficult to assign globally agreeable, consistent, and unique labels for each agent within an open system.

In the next section, we develop a novel solution to the problem of rumour propagation within decentralised reputation systems. Our solution is based on an approach developed within the area of target tracking and data fusion [9]. It avoids the need to uniquely identify an agent, it allows agents to restrict the number of other agents who they reveal their private estimates to, and yet it still allows information to propagate throughout the network.

6.1 Private and Shared Information

Our solution to rumour propagation within decentralised reputation systems introduces the notion of *private* information that an agent knows it has not communicated to any other agent, and *shared* information that has been communicated to, or received from, another agent. Thus, the agent can decompose its contract outcome vector, $\mathcal{N}$, into two vectors, a private one, $\mathcal{N}_p$, that has not been communicated to another agent, and a shared one, $\mathcal{N}_s$, that has been shared with, or received from, another agent:

$$\mathcal{N} = \mathcal{N}_p + \mathcal{N}_s \tag{19}$$

Now, whenever an agent communicates reputation, it communicates both its private and shared vectors separately. Both the originating and receiving agents then update their two vectors appropriately. To understand this, consider the case that agent a_α sends its private and shared contract outcome vectors, $\mathcal{N}_p^\alpha$ and $\mathcal{N}_s^\alpha$, to agent a_β that itself has private and shared contract outcomes $\mathcal{N}_p^\beta$ and $\mathcal{N}_s^\beta$. Each agent updates its vectors of contract outcomes according to the following procedure:

- **Originating Agent:** Once the originating agent has sent both its shared and private contract outcome vectors to another agent, its private information is no longer private. Thus, it must remove the contract outcomes that were in its private vector, and add them into its shared vector:

$$\mathcal{N}_s^\alpha \leftarrow \mathcal{N}_s^\alpha + \mathcal{N}_p^\alpha$$
$$\mathcal{N}_p^\alpha \leftarrow \emptyset.$$

- **Receiving Agent:** The goal of the receiving agent is to accumulate the largest number contract outcomes (since this will result in the most precise estimate) without including shared information from both itself and the other agent (since this may result in double counting of contract outcomes). It has two choices depending on the total number of contract outcomes[6] within its own shared vector, N_s^β, and within that of the originating agent, N_s^α. Thus, it updates its vector according to the procedure below:

  - $N_s^\beta > N_s^\alpha$: If the receiving agent's shared vector represents a greater number of contract outcomes than that of the shared vector of the originating agent, then the agent combines its shared vector with the private vector of the originating agent:

$$\mathcal{N}_s^\beta \leftarrow \mathcal{N}_s^\beta + \mathcal{N}_p^\alpha$$
$$\mathcal{N}_p^\beta \quad \text{unchanged.}$$

  - $N_s^\beta < N_s^\alpha$: Alternatively if the receiving agent's shared vector represents a smaller number contract outcomes than that of the shared vector of the originating agent, then the receiving agent discards its own shared vector and forms a new one from both the private and shared vectors of the originating agent:

$$\mathcal{N}_s^\beta \leftarrow \mathcal{N}_s^\alpha + \mathcal{N}_p^\alpha$$
$$\mathcal{N}_p^\beta \quad \text{unchanged.}$$

  In the case that $N_s^\beta = N_s^\alpha$ then either option is appropriate. Once the receiving agent has updated its sets, it uses the contract outcomes within both to form its trust estimate. If agents receive several vectors simultaneously, this approach generalises to the receiving agent using the largest shared vector, and the private vectors of itself and all the originating agents to form its new shared vector.

This procedure has a number of attractive properties. Firstly, since contract outcomes in an agent's shared vector are never combined with those in the shared vector of another agent, outcomes that originated from the same agent are never combined together, and thus, rumour propagation is completely avoided. However, since the receiving agent may discard its own shared vector, and adopt the shared vector of the originating agent, information is still propagated around the network. Moreover, since contract outcomes are aggregated together within the private and shared vectors, the message size is constant and does not increase as the number of interactions increases. Finally, an agent only communicates its own private contract outcomes to its immediate neighbours. If this agent subsequently passes it on, it does so as unidentifiable aggregated information within its shared information. Thus, an agent may limit the number of agents with which it is willing to reveal identifiable contract outcomes, and yet these contract outcomes can still propagate within the network, and thus, improve estimates of other agents. Next, we demonstrate empirically that these properties can indeed be realised in practise.

[6] Note that this may be calculated from $N = n_{00}^{ab}+n_{01}^{ab}+n_{10}^{ab}+n_{11}^{ab}$.

6.2 Empirical Comparison

In order to evaluate the effectiveness of this procedure we simulated random networks consisting of ten agents. Each agent has some direct experience of interacting with a supplier (as described in section 4.3). At each iteration of the simulation, it interacts with its immediate neighbours and exchanges reputation reports through the sufficient statistics of their Dirichlet distributions. We compare our solution to two of the most obvious decentralised alternatives:

- **Private and Shared Information:** The agents follow the procedure described in the previous section. That is, they maintain separate private and shared vectors of contract outcomes, and at each iteration they communicate both these vectors to their immediate neighbours.
- **Rumour Propagation:** The agents do not differentiate between private and shared contract outcomes. At the first iteration they communicate all of the contract outcomes that constitute their direct experience. In subsequent iterations, they propagate contract outcomes that they receive from any of the neighbours, to all their other immediate neighbours.
- **Private Information Only:** The agents only communicate the contract outcomes that constitute their direct experience.

In all cases, at each iteration, the agents use the Dirichlet distribution in order to calculate their trust estimates. We compare these three decentralised approaches to a centralised reputation system:

- **Centralised Reputation:** All the agents pass their direct experience to a centralised reputation system that aggregates them together, and passes this estimate back to each agent.

This centralised solution makes the most effective use of information available in the network. However, most real world problems demand decentralised solutions due to scalability, modularity and communication concerns. Thus, this centralised solution is included since it represents the optimal case, and allows us to benchmark our decentralised solution.

The results of these comparisons are shown in figure 4. Here we show the sum of the information content of each agent's covariance matrix (calculated as discussed earlier in section 4.3), for each of these four different approaches. We first note that where private information only is communicated, there is no change in information after the first iteration. Once each agent has received the direct experience of its immediate neighbours, no further increase in information can be achieved. This represents the minimum communication, and it exhibits the lowest total information of the four cases. Next, we note that in the case of rumour propagation, the information content increases continually, and rapidly exceeds the centralised reputation result. The fact that the rumour propagation case incorrectly exceeds this limit, indicates that it is continuously counting the same contract outcomes as they cycle around the network, in the belief that they are independent events. Finally, we note that using private and shared information represents a compromise between the private information only case and the centralised reputation case. Information is still allowed to propagate around the network, however rumours are eliminated.

As before, we also plot a single instance of the trust estimates from one agent (i.e. $\hat{p}(X)$ and $\text{Cov}(p(X))$) as a set of ellipses on a

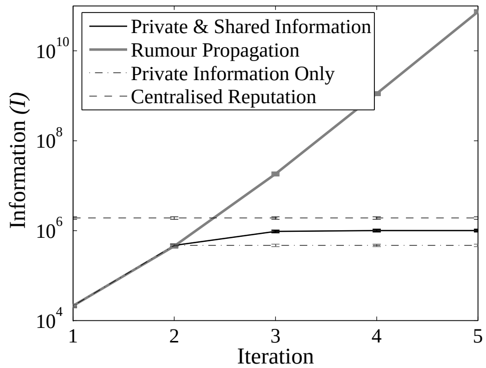

Figure 4: Sum of information over all agents as a function of the communication iteration.

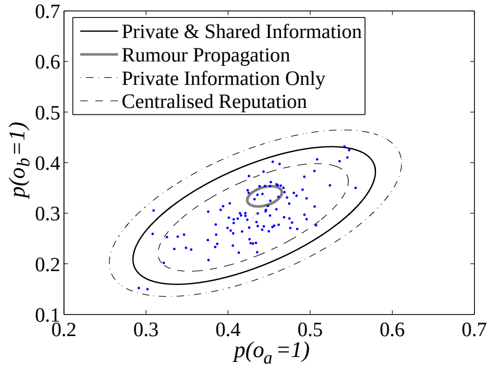

Figure 5: Instances of $\hat{p}(X)$ and Cov$(p(X))$ plotted as second standard error ellipses after 5 communication iterations.

two-dimensional plane (along with samples from the true distribution). As expected, the centralised reputation system achieves the best estimate of the true distribution, since it uses the direct experience of all agents. The private information only case shows the largest ellipse since it propagates the least information around the network. The rumour propagation case shows the smallest ellipse, but it is inconsistent with the actual distribution $p(X)$. Thus, propagating rumours around the network and double counting contract outcomes in the belief that they are independent events, results in an overconfident estimate. However, we note that our solution, using separate vectors of private and shared information, allows us to propagate more information than the private information only case, but we completely avoid the problems of rumour propagation.

Finally, we consider the effect that this has on the agents' calculation of the expected utility of the contract. We assume the same utility function as used in section 4.3 (i.e. $u(o_a = 1) = 6$ and $u(o_b = 1) = 2$), and in table 1 we present the estimate of the expected utility, and its standard deviation calculated for all four cases by a single agent at iteration five (after communication has ceased to have any further effect for all methods other than rumour propagation). We note that the rumour propagation case is clearly inconsistent with the centralised reputation system, since its standard deviation is too small and does not reflect the true uncertainty in the expected utility, given the contract outcomes. However, we observe that our solution represents the closest case to the centralised reputation system, and thus succeeds in propagating information throughout the network, whilst also avoiding bias and overconfidence. The exact difference between it and the centralised reputation system depends upon the topology of the network, and the history of exchanges that take place within it.

Method	$E[E[U]] \pm \sqrt{\text{Var}(E[U])}$
Private and Shared Information	3.18 ± 0.54
Rumour Propagation	3.33 ± 0.07
Private Information Only	3.20 ± 0.65
Centralised Reputation	3.17 ± 0.42

Table 1: Estimated expected utility and its standard error as calculated by a single agent after 5 communication iterations.

7. CONCLUSIONS

In this paper we addressed the need for a principled probabilistic model of computational trust that deals with contracts that have multiple correlated dimensions. Our starting point was an agent estimating the expected utility of a contract, and we showed that this leads to a model of computational trust that uses the Dirichlet distribution to calculate a trust estimate from the direct experience of an agent. We then showed how agents may use the sufficient statistics of this Dirichlet distribution to represent and communicate reputation within a decentralised reputation system, and we presented a solution to rumour propagation within these systems.

Our future work in this area is to extend the exchange of reputation to the case where contracts are not homogeneous. That is, not all agents observe the same contract dimensions. This is a challenging extension, since in this case, the sufficient statistics of the Dirichlet distribution can not be used directly. However, by addressing this challenge, we hope to be able to apply these techniques to a setting in which a suppliers provides a range of services whose failures are correlated, and agents only have direct experiences of different subsets of these services.

8. ACKNOWLEDGEMENTS

This research was undertaken as part of the ALADDIN (Autonomous Learning Agents for Decentralised Data and Information Networks) project and is jointly funded by a BAE Systems and EPSRC strategic partnership (EP/C548051/1).

9. REFERENCES

[1] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan. *Estimation with Applications to Tracking and Navigation*. Wiley Interscience, 2001.

[2] C. Boutilier. The foundations of expected expected utility. In *Proc. of the 4th Int. Joint Conf. on on Artificial Intelligence*, pages 285–290, Acapulco, Mexico, 2003.

[3] M. Evans, N. Hastings, and B. Peacock. *Statistical Distributions*. John Wiley & Sons, Inc., 1993.

[4] N. Griffiths. Task delegation using experience-based multi-dimensional trust. In *Proc. of the 4th Int. Joint Conf. on Autonomous Agents and Multiagent Systems*, pages 489–496, New York, USA, 2005.

[5] N. Gukrai, D. DeAngelis, K. K. Fullam, and K. S. Barber. Modelling multi-dimensional trust. In *Proc. of the 9th Int. Workshop on Trust in Agent Systems*, Hakodate, Japan, 2006.

[6] A. Jøsang and R. Ismail. The beta reputation system. In *Proc. of the 15th Bled Electronic Commerce Conf.*, pages 324–337, Bled, Slovenia, 2002.

[7] E. M. Maximilien and M. P. Singh. Agent-based trust model involving multiple qualities. In *Proc. of the 4th Int. Joint Conf. on Autonomous Agents and Multiagent Systems*, pages 519–526, Utrecht, The Netherlands, 2005.

[8] S. D. Ramchurn, D. Hunyh, and N. R. Jennings. Trust in multi-agent systems. *Knowledge Engineering Review*, 19(1):1–25, 2004.

[9] S. Reece and S. Roberts. Robust, low-bandwidth, multi-vehicle mapping. In *Proc. of the 8th Int. Conf. on Information Fusion*, Philadelphia, USA, 2005.

[10] J. Sabater and C. Sierra. REGRET: A reputation model for gregarious societies. In *Proc. of the 4th Workshop on Deception, Fraud and Trust in Agent Societies*, pages 61–69, Montreal, Canada, 2001.

[11] W. T. L. Teacy, J. Patel, N. R. Jennings, and M. Luck. TRAVOS: Trust and reputation in the context of inaccurate information sources. *Autonomous Agents and Multi-Agent Systems*, 12(2):183–198, 2006.

[12] S. Utete. *Network Management in Decentralised Sensing Systems*. PhD thesis, University of Oxford, UK, 1994.