GMD Report
44

GMD –
Forschungszentrum
Informationstechnik
GmbH

Peter Fankhauser, Thomas Tesch

Agents, a Broker, and Lies

November 1998

© GMD 1998

GMD –
Forschungszentrum Informationstechnik GmbH
Schloß Birlinghoven
D-53754 Sankt Augustin
Germany
Telefon +49 -2241 -14 -0
Telefax +49 -2241 -14 -2618
http://www.gmd.de

In der Reihe GMD Report werden Forschungs- und Entwicklungsergebnisse aus der GMD zum wissenschaftlichen, nichtkommerziellen Gebrauch veröffentlicht. Jegliche Inhaltsänderung des Dokuments sowie die entgeltliche Weitergabe sind verboten.

The purpose of the GMD Report is the dissemination of research work for scientific non-commercial use. The commercial distribution of this document is prohibited, as is any modification of its content.

Anschrift der Verfasser/Address of the authors:
Dr. Peter Fankhauser
Thomas Tesch
Institut für Integrierte Publikations- und Informationssysteme
GMD – Forschungszentrum Informationstechnik GmbH
Dolivostraße 15
D-64293 Darmstadt
E-mail: {fankhaus,tesch}@darmstadt.gmd.de
http://www.darmstadt.de/~{fankhaus,tesch}

ISSN 1435-2702

Abstract

Virtual enterprises need reliable and efficient cooperation mechanisms to carry out transactions between autonomous agents with conflicting interests. Available cooperation mechanisms either use bilateral multi-step negotiation or auctioning. Negotiations encourage agents to reason about the interests of their opponents. Thus, negotiations suffer from counterspeculations. Auctions apply to asymmetric trading only; they either favor the auctioneer or the bidders. Both mechanisms do not promote agents to tell the truth. Therefore, we propose to use a trustbroker to mediate between the agents. We introduce three symmetric, negotiation free one-step protocols to carry out a sequence of decisions for agents with possibly conflicting interests. The protocols achieve substantially better overall benefit than random or hostile selection, and they avoid lies. We analyze the protocols with respect to informed vs. uninformed lies, and with respect to beneficial vs. malevolent lies, and show that agents are best off to know and announce their true interests. Analytical proofs and simulations substantiate our results.

Keywords: Agents, Broker, Lies, Cooperation, Transaction, Negotiation, Mediation

Zusammenfassung

Virtuelle Organisationen benötigen zuverlässige und effiziente Kooperationsmechanismen, um Transaktionen zwischen autonomen Agenten mit konfligierenden Interessen durchzuführen. Bestehende Kooperationsmechanismen basieren auf bilateraler Verhandlung oder Auktionen. Verhandlungsprotokolle ermutigen Agenten, die Interessen ihrer Gegenspieler herauszufinden, was zu wechselseitigen Spekulationen führt. Auktionen sind nur auf asymmetrische Situationen anwendbar. Sie bevorzugen entweder den Auktionator oder die Bieter. Beide Ansätze motivieren die Agenten nicht, die Wahrheit zu sagen. Daher schlagen wir die Verwendung eines vertrauenswürdigen Maklers vor, der zwischen den Agenten vermittelt. Wir führen drei symmetrische, verhandlungsfreie Ein-Schritt-Protokolle ein, die eine Folge von Entscheidungen für Agenten mit konfligierenden Interessen durchführen. Die Protokolle erreichen einen wesentlich höheren Nutzen als zufällige oder feindselige Entscheidungen, und sie vermeiden Lügen. Wir analysieren die Protokolle in Hinblick auf informierte vs. uninformierte Lügen und in Hinblick auf vorteilhafte vs. böswillige Lügen, und zeigen, daß die Agenten den höchsten Nutzen mit der Wahrheit erreichen. Analytische Beweise und Simulationen untermauern unsere Resultate.

Schlüsselworte: Agenten, Makler, Lügen, Kooperation, Transaktion, Verhandlung, Vermittlung

1 Introduction

Virtual enterprises need reliable and efficient cooperation mechanisms to carry out transactions between autonomous agents. In such scenarios, transactions are often ad-hoc and thus can only be executed with a dynamically determined plan. Furthermore, the agents are self-interested which prevails agreement on a common plan.

A key problem in realizing cooperation in such settings is to have reliable means for choosing a transaction from a number of alternatives in such a way that it compromises the conflicting interests of the participating agents [TA98]. Existing cooperation mechanisms can be classified into bilateral multi-step negotiations [ZR89, RZ94] and auctioning [Var95].

Bilateral negotiations consider symmetric situations where two agents try to agree on a common plan. Many of the approaches to negotiation stem from research in game theory [NM47, FT91], which traditionally assumes public information among multiple agents [ZR89]. This does not apply to agents acting in competitive and hostile environments [ZS96]. The negotiation approach presented in [SFJ97, NJFM96, JFN$^+$96] for the domain of business process management considers privacy of information. However, it requires agents to reason about the interests and strategies of their opponents, and thus, suffers from counterspeculations. More gravely, an agent can be better off knowing the interests of its opponents and pretend different interests to reach its true objective.

Another line of research for matching the interests of autonomous agents are market based mechanisms like auctioning. Auctions consider asymmetric situations where an auctioneer wants to sell an item for the highest possible price, and bidders want to acquire the item for the lowest possible price. In [GM98a, GMM98, GM98b] it has been pointed out that auction protocols either favor the bidder or the seller, forcing buyers and sellers into price-wars and provoking various manipulations like fake bidders (shills), sellers acting as auctioneers, or coalition formation by bidders [SL95]. Also for the widely used Vickrey auction [Vic61] it has been shown that bidding the truth is not always the dominant strategy for an agent [San95].

In order to overcome the shortcomings of bilateral negotiations and auctions, we propose to mediate conflicting interests with a neutral third instance, a trustbroker. The agents report their interests to the trustbroker which then selects an optimal and fair compromise. To avoid manipulations, the trustbroker employs a selection scheme which ensures that agents are best off to know and announce their true interests.

Having an independent intermediary mediating the interests of two opponents is a common setting in real life business interactions, often implemented by a commonly chosen lawyer or mediator [Rai82]. For electronic business to business cooperations, mediation is even more important. Apart from mediating conflicting interests among agents which is the focus of this paper, brokers are also developed for finding agents with matching interests [Fon97], integrating heterogeneous information sources [KH96, JBB$^+$97], and realizing authentication and security [Tyg98].

The remainder of the paper is organized as follows. In Section 2, we introduce our basic setting, and point out that optimal compromises give rise to simple uninformed lies. Therefore, we present two protocols which almost achieve the optimal compromise and avoid uninformed lies. To avoid also informed lies, we develop another protocol in Section 3 which still achieves significantly better benefit than random conflict resolution. We illustrate the behavior of the protocols by simulations in Section 4 and summarize and discuss future work in Section 5.

2 Deterministic Protocols

2.1 The Setting

In this paper we investigate the following setting. Two agents want to carry out an open-ended sequence of decisions. The agents rank the alternatives of each decision differently, possibly with conflicts.

We give a simple example of a single decision. One agent (a travel agency) wants to sell a flight ticket together with a hotel voucher. It prefers to sell them as a package, so it has a higher utility for selling {flight, hotel} than for selling either {flight} or {hotel} alone. The other agent (the client) prefers to book the flight only, so it assigns rather high utility to {flight}, but smaller utility to {flight, hotel}, and probably zero utility to {hotel}.

Like in the above example, we assume that all actions are complete exchanges—booking involves reservation and advance payment—and that actions can consist of a sequence of atomic actions.

Handling such situations today is usually unbalanced and inflexible. The travel agency typically only offers the package. It assigns the maximal utility to the package, and the client either accepts these conditions, or looks for a more adequate offer. However, the travel agency may also be interested in filling up its flight contingent, and the client may as well accept the full package, possibly planning to resell the hotel voucher. Thus the agents could match, if they are able to find a compromise.

Our goal is to find and execute such compromises. To this end we introduce a trustbroker. The trustbroker is a trusted entity, to which agents can announce their utilities for a particular set of actions. The trustbroker selects the action where the sum of utilities is maximal, and it can force the agents to execute the selected action.

More formally, for each decision at timestep $t \in \mathbb{N}$, among n_t alternative actions $a_{i,t}, i = 0 \ldots n_t - 1$, with utilities $u_{i,t}$ for one agent and utilities $v_{i,t}$ for the other agent, the trustbroker selects the action where $(u_{i,t} + v_{i,t})$ is maximal. To make the utilities comparable we normalize them such that $\sum u_{i,t} = 1$ and $\sum v_{i,t} = 1$. When regarding only one decision we will omit the index t.

2.2 The Optimum and Greedy Lies

With true announced utilities which are distributed equally for both agents, the above selection scheme amounts in a fair and optimal overall benefit for the agents. To assess the average utility per step gained by one agent we need to make assumptions about the distribution of utilities.

Let us consider the binary case first. For two actions with utilities $u, 1-u$, and $v, 1-v$, and u and v uniformly distributed between 0 and 1 the average utility gained by each agent can be calculated as follows. Averaging over all v, the probability $p(u+v > 1-u+1-v)$ is u. Thus agent 1 with utilities u and $1-u$ gets $p(u+v > 2-u-v)u+(1-p(u+v > 2-u-v))(1-u) = u^2+(1-u)^2$. Averaging over all utilities u amounts in $\int_0^1 u^2 + (1-u)^2 \, \mathrm{d}u = \frac{2}{3}$.

For higher n this distribution can be naturally generalized as follows. Let x_i, $0 < i < n$, be $n-1$ values, all uniformly and independently distributed between 0 and 1, sorted in ascending order such that $x_i \leq x_{i+1}$, and let $x_0 = 0$ and $x_n = 1$. Then each utility $u_i = x_{i+1} - x_i$.

Thereby each u_i is the minimum of $(n-1)$ x_i's. The resulting density of each u_i is $f(u_i, n) = (n-1)(1-u_i)^{n-2}$ which is a probability mass function. The mean is $\frac{1}{n}$. With higher n, the overall utility is distributed among more alternatives, and the likelihood of a conflict increases. Thus the overall utility gained decreases. For n between 2 and 15, this highest average gain is depicted in Figure 2 by the maximum sum selection (a).

The optimum overall benefit can only be achieved with honest agents which announce their true utilities. A self-interested agent achieves a better benefit by lying its better utilities up, and decreasing the other utilities accordingly. In the extreme, it can lie its better utility up to 1, and thereby virtually ensure the selection of the better action. For the binary case this leads to an average utility of $\frac{3}{4}$ for the greedy agent (as opposed to $\frac{2}{3}$), whereas the honest agent gets only $\frac{1}{2}$. If both agents adopt this greedy strategy, the trustbroker has to perform a random draw for conflicting actions, leading to an average utility of $\frac{5}{8}$ for both. So both agents lose. With higher n, this greedy strategy performs even worse, because the probability of a conflict increases drastically.

2.3 Deterministic Selection with Preferences

One way to avoid greedy lies is as follows. Rather than announcing their utilities, the agents need only rank the actions according to their preferences, assigning a unique rank $0 \le r_i < n$ to each action a_i, such that $r_i < r_j$ if $u_i \le u_j$. The broker then selects the action with the minimum sum of ranks. This approach does not result in the best overall utility; for the binary case it just amounts in the mutual greedy lie described above which gives $\frac{5}{8}$. Also for higher n it performs slightly worse than the optimal selection as illustrated by preference selection (b) in Figure 2. However, it performs significantly better than an approach where both agents greedily lie their best utility to 1. This is because the preference selection finds a compromise for conflicting utilities, whereas the mutual greedy lie leads to a random draw for conflicts.

The preference selection avoids greedy beneficial lies; every false ranking damages on average the benefit of the lying agent more than the benefit of the honest agent. However, when one agent knows the preferences of the other agent, it can announce false utilities to get a better benefit. Consider the following simple example. We have three actions with $u_0 = 0.1, u_1 = 0.3, u_2 = 0.6$, and $v_0 = 0.51, v_1 = 0.49, v_2 = 0$. The first agent thus ranks the actions with $[3, 2, 1]$ and the second agent ranks them with $[1, 2, 3]$. In this case the sum of the ranks is $[4, 4, 4]$ resulting in random selection. Thus, both agents can only expect to get $\frac{1}{3}$. Assume agent 2 knows the preferences of its opponent and thus knows that there is a complete conflict. Lying its ranks to $[2, 1, 3]$ leads to rank-sums $[5, 3, 4]$; thus the second action is selected. With this selection agent 2 gets 0.49 and agent 1 gets only 0.3.

2.4 Deterministic Selection with Weights

Another way to avoid greedy lies is to attach a cost to each utility announcement. This can be accomplished as follows. Both agents get an initially equal budget. An agent's budget determines its influence on the decision of the trustbroker. Furthermore, at each step the trustbroker adjusts the individual budgets according to the agents' expected gain. When the expected gain of the first agent is higher than the gain of the second, the first agent's budget is increased, otherwise it

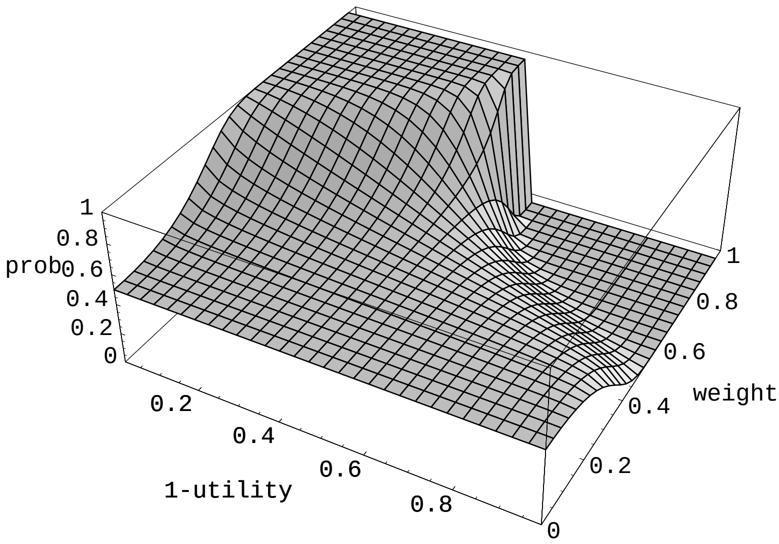

Figure 1: Selection probability for deterministic selection with weights

is decreased.

We represent budgets by weights w for agent 1 and $1 - w$ for agent 2, with $0 \leq w \leq 1$. On this basis, we generalize the maximum sum selection scheme by choosing the action a_i where $u_i^{\frac{w}{1-w}} v_i^{\frac{1-w}{w}}$ is maximal.

For the binary case, averaging over all v yields for weight w and an action with utility u the selection probability $p(u^{\frac{w}{1-w}} v^{\frac{1-w}{w}} > (1-u)^{\frac{w}{1-w}} (1-v)^{\frac{1-w}{w}})$, which is depicted in Figure 1. For $w = \frac{1}{2}$ this strategy is equivalent to determining the maximum sum, whereas for $w > \frac{1}{2}$ it drastically increases the influence of utilities u_i. However, other than with using a maximum weighted sum, an agent cannot completely determine the outcome for a selection, unless it has weight 1 or it announces an action with utility 1. To avoid the latter we restrict all announced utilities to the interval $[\frac{\epsilon}{n}, 1 - \frac{\epsilon}{n}]$, with, e.g., $\epsilon = 0.1$.

To guard against greedy lies, the trustbroker adjusts the weights as follows:

$$w_{t+1} = (1 - w_t)(\sigma^2(u_{i,t}) - \sigma^2_{max}) - w_t(\sigma^2(v_{i,t}) - \sigma^2_{max})$$

$\sigma^2(u_{i,t})$ is the standard deviation of $u_{i,t}$ at timestep t

$\sigma^2_{max} = \dfrac{n-1}{n^2}$ is the maximal standard deviation

With this adjustment the weight decreases for a high current weight and a high standard deviation of utilities.

This protocol gives almost the optimum overall benefit when both agents tell the truth. Because the weights stay close to $\frac{1}{2}$, in almost all cases the action with the maximal sum of utilities is chosen.

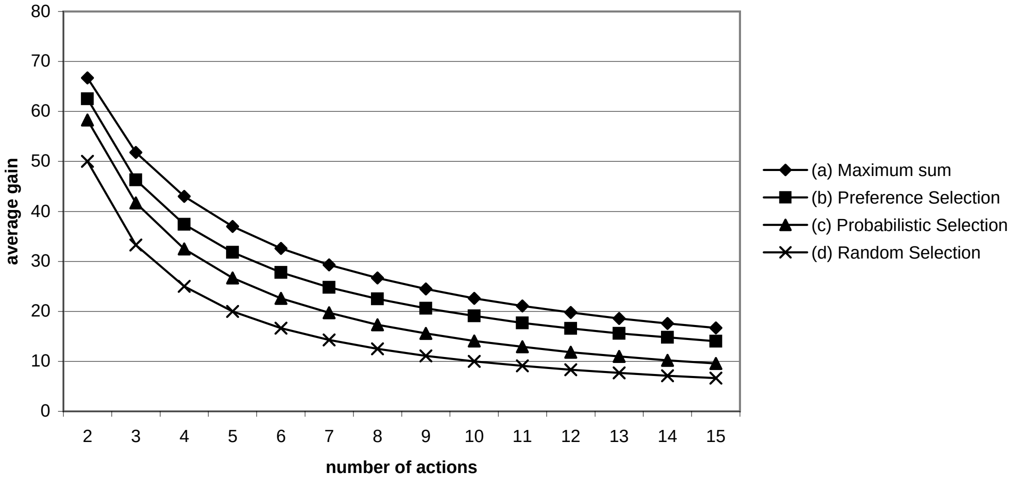

Figure 2: Comparison of different selection schemes with two honest agents

Our simulations (Section 4) indicate that this protocol indeed avoids greedy lies. But it also avoids cautious lies, where one agent artificially lowers the standard deviation of its utilities in order to increase its future gain, because a high influence is immediately punished in the next step. However, like with the deterministic selection with preferences the protocol does not avoid informed lies, as exemplified by the beneficial informed strategy in Section 4.

3 Probabilistic Protocol

3.1 Why Nondeterminism and How

When the selection of an action is completely determined by the announced utilities, and possibly by the weights, an agent who knows the utilities of its opponent can always lie its own utility to enforce the selection of the preferred action, or also to increase its future influence. Thus there can be informed lies.

To overcome this problem, we need to introduce some sort of nondeterminism. If an agent can not completely foresee which action is chosen, even if it has complete knowledge about the involved utilities and weights, it can also not rely on the effect of an informed lie with certainty. As we will show in Sections 3.2 and 3.3, this uncertainty suffices to avoid informed beneficial lies.

We accomplish nondeterminism with the following selection and weight adjustment scheme. With action utilities $u_{i,t}, v_{i,t}$, and weights w_t, $1 - w_t$ as defined above, the trustbroker selects action a_i with probability $p_{i,t}$ and adjusts weights to w_{t+1} on the basis of the expected gain $Gain()$ of each agent:

$$p_{i,t} = w_t u_{i,t} + (1 - w_t) v_{i,t}$$
$$w_{t+1} = w_t - Gain(u_{0,t}, \ldots, u_{n-1,t}, w_t) + Gain(v_{0,t}, \ldots, v_{n-1,t}, 1 - w_t)$$

The larger the weighted sum of utilities for an action is, the more likely it will be selected.

Furthermore, like with the weighted deterministic selection, an agent who can gain more than the other agent in one step will have less influence on the selection in the future.

To arrive at a concrete weight adjustment scheme, we need to determine the average gain of an agent depending on its utilities and weight. This can be calculated on the basis of the distribution of n utilities introduced in Section 2.2. The probability $p(a_i)$ that an action a_i with utility u_i is selected, when the corresponding utility v_t of agent 2 is distributed with density $f(v_i, n)$ is:

$$p(a_i) = \int_0^1 (wu_i + (1-w)v_i f(v_i, n))\ \mathrm{d}v_i = \frac{1}{n} + w(u_i - \frac{1}{n}) \text{ for all } n > 1$$

The expected gain in one step is thus:

$$Gain(u_0, \ldots u_{n-1}, w) = \sum_{i=0}^{n-1} p(a_i)u_i = \frac{1}{n} + w\sum_{i=0}^{n-1}(u_i - \frac{1}{n})u_i$$

As can be easily verified, $Gain()$ takes the minimal value of $\frac{1}{n}$ for $u_0 = \ldots = u_{n-1} = \frac{1}{n}$, and the maximal value of $\frac{1}{n} + w\frac{n-1}{n}$ for one $u_j = 1$ and all other $u_i = 0$. Thus, the more an agent prefers a particular action, the more it can expect to gain.

Because this protocol also selects suboptimal actions with a certain probability, the agents will accumulate less overall benefit than with the deterministic protocols. However, because better actions are selected more likely, the agents will still get a significantly better benefit than a completely random selection, which would just gain $\frac{1}{n}$ per step. Let us consider the ratio of random gain vs. probabilistic gain in more detail. Based on $Gain()$, the average gain $AvgGain(w, n)$ in one step can be calculated by:

$$AvgGain(w, n) = n\int_0^1 f(u_i, w)p(a_i)u_i\ \mathrm{d}u_i = \frac{1 + n - w + nw}{n(1+n)}$$

The weights oscillate around $\frac{1}{2}$. So for $w = \frac{1}{2}$ the average gain is $\frac{1+3n}{2n(1+n)}$ (see also probabilistic selection (c) in Figure 2). For large n the ratio thus converges to

$$\lim_{n\to\infty} \frac{1}{n}\frac{2n(1+n)}{1+3n} = \frac{2}{3}$$

3.2 Informed Lies

Our main motivation for introducing a probabilistic selection scheme is to avoid informed lies. A lie announces false utilities to either increase the immediate gain, or to increase the future influence, and thereby increase the future gain. An informed lie can occur, if an agent needs to know the utilities of the other agent in order to calculate the effect of the lie. But it can be shown that the effect of any possible lie can be calculated without knowing the utilities of the other agent.

Theorem 1 *The effect of a lie can be calculated without knowing of the utilities of the other agent.*

Proof 1 Let us assume an agent with weight w and true utilities u_i, lies each utility by some δ_i, with $\sum_{i=0}^{n-1} \delta_i = 0$ and $-u_i \leq \delta_i \leq 1 - u_i$, whereas the other agent announces its true utilities v_i.

(a) The probability that an action with announced utilities $u_i + \delta_i$ and v_i is selected is $p_i^{\delta_i} = w(u_i + \delta_i) + (1 - w)v_i$. With the true announced utility u_i it is $p_i = wu_i + (1 - w)v_i$. The effect of the lie on the immediate gain is the difference between the utility gained with lying and the utility gained without lying.

$$\sum_{i=0}^{n-1} p_i^{\delta_i} u_i - \sum_{i=0}^{n-1} p_i u_i = w \sum_{i=0}^{n-1} \delta_i u_i$$

The utilities v_i of the opponent cancel out; thus the effect of the lie does not depend on them.

(b) Similarly, the effect of the lie on the future influence is the difference between the next weight with a lie and the next weight without a lie.

$$\begin{aligned}
& w - Gain(u_0 + \delta_0, \ldots, u_{n-1} + \delta_{n-1}, w) + Gain(v_0, \ldots, v_{n-1}, 1 - w) - \\
& w + Gain(u_0, \ldots, u_{n-1}, w) - Gain(v_0, \ldots, v_{n-1}, 1 - w) \\
& = w \sum_{i=0}^{n-1} \delta_i (\frac{1}{n} - 2u_i - \delta_i)
\end{aligned}$$

Thus also the effect of the lie on the future influence does not depend on the utilities of the opponent. □

3.3 Beneficial and Malevolent Lies

In the previous section we have shown that the effect of a lie on the immediate gain and on the future influence does not depend on the utilities of the opponent. Of course a lie does have an effect. An agent can still optimize its immediate gain by lying the utility of its best action up to 1. Conversely, an agent can optimize its future influence, and thereby its future gain by lying all its utilities down to $\frac{1}{n}$. These two optimizations are in conflict. Lying greedily gains immediately, but costs future influence and vice versa. To avoid benefits from any sort of a lie, we have to guarantee that for all possible lies:

$$gainnow(lie) + gainfuture(lie) \leq gainnow(truth) + gainfuture(truth)$$

Theorem 2 *For all utilities u_i,v_i, all weights w_0, and all lying deltas δ_i the following holds. Let $gain_t(\delta)$ denote the gain in step t with a lie by δ in the first step. Then*

$$\sum_{t=0}^{\infty} gain_t(\delta) - gain_t(0) \leq 0 \textit{ for all } \delta$$

Proof 2 To avoid notational noise, we restrict ourselves to an open-ended sequence of binary decisions, with utilities u_t, $1 - u_t$, and v_t, $1 - v_t$.

(a) the effect of a lie now can be determined from step (a) in Proof 1:

$$gain_0(\delta) - gain_0(0) = (2u_0 - 1)\delta w$$

(b) the effect of this lie on the future gain requires a bit more complication, because we need to analyze not only the immediate next step, but all future steps. For this purpose we introduce the auxiliary function $wchange_t$ which describes the weight change at step t, caused by an initial lie with δ.
$wchange_0 = 2w\delta(1 - 2u_0 - \delta)$, is the change in weight by δ in the next step (recall Proof 1, step (b)).
$wchange_{t+1} = \frac{2}{3} wchange_t$, is the average change in weight depending on the change in weight in the previous step. This can be derived by $\int_0^1 (\int_0^1 w - Gain(u, 1-u, w) + Gain(v, 1-v, 1-w)\ \mathrm{d}v)\ \mathrm{d}u = \frac{1}{6} + \frac{2w}{3}$, which gives the weight change averaged over all u and v in a step. We thus arrive at:

$$wchange_t = 2w\delta(1 - 2u_0 - \delta)(\frac{2}{3})^t$$

Consequently, we arrive at the following figures:

$$\begin{aligned}
&gain_0(\delta) - gain_0(0) = (2u_0 - 1)\delta w \\
&gain_t(\delta) - gain_t(0) = AvgGain(wchange_{t-1}, 2) = \frac{1}{3} w\delta(1 - 2u_0 - \delta)(\frac{2}{3})^t \\
&\sum_{t=0}^{\infty} gain_t(\delta) - gain_t(0) = \delta w(2u_0 - 1 + \frac{1}{3}(1 - 2u_0 - \delta)\sum_{t=1}^{\infty}(\frac{2}{3})^t) = \\
&= -\delta^2 w \leq 0 \text{ for all } \delta, \text{ and all } w \geq 0
\end{aligned}$$

Thus there can be no beneficial lies. □

Note that in computing the average gain of all future steps $t > 0$ we have assumed that both agents do not lie. However, if a lie does not already gain in one step, then also two subsequent lies can not gain, provided that they are followed by an infinite number of steps. In practice, the effect of any greedy lie is compensated by just two or three further steps. And a cautious lie never gains enough future influence such that it can be compensated with an average next step. Nevertheless, if an agent can freely leave a trustbroker and refuse any further cooperation, when it has arrived at a small weight, it can of course lie beneficially. The development of appropriate authentication and risk management mechanisms for the trustbroker is subject for future work.

Beneficial lies consider only the effect of a lie on the utility of the agent who lies. A malevolent agent could have a different policy, which aims at maximizing the difference between the average gain by the malevolent agent and the gain of its opponent, even if that means that it damages its own gain also. Unfortunately, such malevolent lies can not be avoided, but they are never beneficial. One malevolent lie simply maximizes the differences between the own and the opponent's utilities as described in the next section.

4 Simulations

We have performed simulations for two agents and one trustbroker. The number of actions n investigated ranges from 2 to 15, e.g., for each n we have simulated $1,000,000$ steps[1] in which the trustbroker mediates among n actions. In each step, each agent gets assigned its true utilities randomly with the distribution introduced in Section 2.2. To investigate the robustness of selection schemes against different types of lies, we simulate lying agents that announce utilities which deviate from their true utilities. The following strategies are considered:

- *honest (h)*: an honest agent announces in each step its true utilities to the trustbroker.
- *greedy (g)*: a greedy lying agent lies the action with the maximal utility to 1 and all others to 0. The rationale behind this lie is to increase the probability that the best action is selected.
- *semi-greedy (sg)*: this strategy lies only greedily for actions with a relatively high utility, i.e., only if the utility u_i of the best action is greater than the threshold value $\frac{1}{n} + \sigma^2_{max}$, u_i is lied to 1 and the smaller utilities are set to 0. The rationale behind this lie is to be only greedy when it is worth it.
- *malevolent uninformed (mu)*: the objective of a malevolent uninformed lie is to change the own utilities in a way that hurts the opponent more than oneself without any knowledge of the opponent's utilities. Each true u_i is lied to $\frac{1}{2}(u_i + \frac{1}{n})$ which lowers the standard deviation of the announced utilities without changing the preferences. The rationale behind this lie is to group the utilities around $\frac{1}{n}$ in order to increase the weight for budget oriented selection schemes.
- *malevolent informed (mi)*: this strategy is informed, thus, it is based on knowing the true utilities of the opponent. Each true u_i is lied to $\frac{1}{n}(u_i + 1 - v_i)$ which lies the utilities to the difference of the true utilities. The rationale behind this lie is to increase the conflicts to damage the opponent and at the same time to increase the future influence in order to be less damaged than the opponent.
- *beneficial informed (bi)*: the beneficial informed lie also tries to outperform the honest strategy by exploiting the opponent's utilities. For the preference selection, we use the malevolent informed strategy which turns into a beneficial informed lie with this selection scheme. For the deterministic selection, the utility of the best action u_{best} is modified in such a way that its selection is guaranteed with the current weight. The remaining utilities are all set to $\frac{1}{n-1}(1 - u_{bestlie})$ in order to minimize the standard deviation. The rationale behind this lie is to ensure that an agent gets its best action with minimal loss of weight. For the probabilistic selection scheme, an informed beneficial lie makes no sense, because, as shown in Section 3.2, the effect on the own benefit is independent of the opponent's utilities.

[1]Note, this large number is only to get an accurate assessment on the quantitative effect of lies; much shorter sequences in the range of 2-7 steps yield the same qualitative results.

n	*h*	*g*	*h*	*sg*	*h*	*mu*	*h*	*bi*	*h*	*h*
2	62.5	62.5	62.5	62.5	62.5	62.5	54.2	58.3	62.5	62.5
3	47.8	44.9	47.2	45.6	46.3	46.3	35.8	48.1	46.3	46.3
5	32.2	28.9	32.1	29.5	31.8	31.8	21.7	35.3	31.8	31.8
10	19.1	15.3	19.1	15.4	19.1	19.1	11.9	22.4	19.1	19.1
15	14.0	10.4	14.0	10.4	14.0	14.0	08.5	16.9	14.0	14.0

Table 1: $1,000,000$ steps, preference selection

n	*h*	*g*	*h*	*sg*	*h*	*mu*	*h*	*mi*	*h*	*bi*	*h*	*h*
2	74.6	53.0	66.7	65.3	68.7	63.9	58.7	71.4	55.2	71.9	66.3	66.3
3	60.9	35.5	54.6	46.6	55.0	47.7	46.3	53.6	39.8	57.2	51.4	51.3
5	45.6	21.2	43.9	25.4	40.4	32.3	39.5	31.9	28.6	40.0	36.4	36.4
10	29.3	10.4	29.2	10.7	25.2	19.0	28.6	13.4	18.9	22.8	22.2	22.2
15	22.1	06.9	22.1	06.9	19.1	13.6	21.9	08.2	14.2	16.3	16.5	16.5

Table 2: $1,000,000$ steps, deterministic selection

In addition to the lies presented, multiple variations have been simulated. However, we limit our presentation to these strategies because they illustrate the strengths and weaknesses of the different selection schemes. Because the maximal gain can only be achieved when both agents report true utilities, and the maximal sum is selected (see Figure 2), strategies that perform better than two honest agents do not exist. Therefore, we present only results for competitions between one honest agent and a lying opponent.

The results for the preference selection are summarized in Table 1. n gives the number of alternatives; utilities range between 0 and 100. The preference selection ranks the utilities as explained in Section 2.3. In each step, the action with the minimal sum of ranks is selected. For two honest agents it performs slightly worse than the maximum sum selection. Greedy and semi-greedy lying agents are punished with increasing n because the actions beyond the first choice become ranked randomly by the selection scheme. But a beneficial informed lie can be accomplished by applying the malevolent informed strategy which maximizes the conflicts between the two agents.

Table 2 presents the results of the deterministic selection scheme with weight adjustment introduced in Section 2.4. The protocol achieves for two honest agents almost the theoretical maximum. It successfully avoids greedy, semi-greedy, and malevolent uninformed lies. The malevolent informed lie is beneficial for $n \leq 3$, but it is successfully avoided by the selection scheme for $n > 3$. These positive properties are neutralized by beneficial informed lies for $n > 2$ which encourages agents to reason about their opponents.

The results of the probabilistic selection scheme are summarized in Table 3. The strategy performs well against greedy lies and semi-greedy lies; the honest agent benefits from the lies of its

n	*h*	*g*	*h*	*sg*	*h*	*mu*	*h*	*mi*	*h*	*bi*	*h*	*h*
2	62.5	56.3	60.3	58.0	53.4	56.6	50.0	55.6	57.5	57.3	58.3	58.3
3	46.6	38.9	45.3	40.3	36.7	40.0	31.8	37.9	39.9	39.8	41.7	41.6
5	31.4	23.7	31.0	24.2	22.7	25.3	18.5	22.5	23.5	23.7	26.7	26.6
10	17.5	11.6	17.5	11.7	11.6	13.3	09.4	10.8	11.1	11.3	14.1	14.1
15	12.2	07.6	12.2	07.6	07.8	09.0	06.3	07.1	07.2	07.3	09.6	09.6

Table 3: $1,000,000$ steps, probabilistic selection

opponents. The malevolent-uninformed strategy is successfully hurting the honest agent more than the liar. This strategy could be counteracted with a malevolent uninformed lie. Repeated mutual counterlies result in a random selection with average gain $\frac{1}{n}$. Also informed malevolent lies can be compensated with an informed malevolent counter lie and run into $\frac{1}{n}$ with repeated counter lies. Furthermore, an informed malevolent lie is almost compensated by an uninformed malevolent lie. Finally, beneficial informed lies are not possible, which has also been shown analytically in Section 3.

Although the probabilistic selection scheme obtains less overall gain for two honest agents in comparison with the other discussed selection schemes, it is the only selection scheme that avoids beneficial informed lies. Thus, even in a hostile environment an agent cannot benefit from knowing its opponent utilities or reasoning about its opponents interests. A mediation based on a deterministic selection scheme still promises increased benefit from spying out the opponent and also contains the risk of being spied out.

In addition, we have experimented with sequences of decisions with variable number of alternatives n, and with finite sequences. The simulations with decreasing or randomly chosen n amount in the average of the presented figures for fixed n. The simulations with finite sequences show that greedy lies for the weighted selection schemes are usually compensated in only 2-3 steps, and at most after 7-8 steps. For the preference selection the sequence length does not make a difference.

5 Conclusion and Future Work

We have introduced several protocols to carry out transactions between self-interested agents mediated by a trustbroker. The protocols have been designed to find compromises, which maximize the overall benefit for the participating agents and avoid lies. These two goals are in conflict. Avoiding lies requires to sometimes select suboptimal actions. Thus, the optimal compromise can not always be achieved. Protocols with deterministic selection achieve relatively good overall benefit, but suffer from informed lies. Protocols with nondeterministic selection are robust against informed beneficial lies, but achieve less overall benefit. However, all presented protocols achieve a balanced benefit, which is significantly better than the benefit of a partially random selection resulting from resolving conflicts without mediation.

Future work will be devoted to the following issues.

We need to analyze the behavior of the protocols for finite decision sequences consisting of partially interdependent decisions. Furthermore, we plan to extend the protocols for dealing with more than two agents.

The role of the trustbroker needs to be extended to insure against agents which drop out of a cooperation. To some extent this can be achieved by appropriate authentication and secure transaction [Bil98, KGM96, Tyg98] mechanisms, and by detecting lying patterns. If these mechanisms fail, the broker should compensate. For this purpose it needs commission mechanisms to accumulate assets from successfully executed transactions, which it can use for compensation.

Finally, the effect of embedding the protocols into mediated matchmaking [Fon97, KH96] needs to be considered. A broker which preferably matches agents with compatible interests will achieve better overall benefit than what can be achieved with the utility distributions used in this paper. On the other hand, matchmaking can be made more efficient, when the execution of a match can rely on the reported utilities, and more flexible, when the actual transaction needs not be determined completely in advance.

Acknowledgments

We would like to thank the OASYS team at IPSI for the inspiring atmosphere enabling this paper. We are indebted to Justus Klingemann for contributing the simple and robust preference selection scheme, to Karl Aberer for contributing the elegant generalization of the binary distribution of utilities to larger n, and to Karl, and to Jürgen Wäsch for questioning everything questionable in a good mood.

References

[Bil98] D. Billard. Transactional services for the internet. In *Proc. of the Int. Workshop on the Web and Databases*, pages 11–17, Valencia, Spain, 1998.

[Fon97] Leonard N. Foner. Yenta: A multi-agent, referral-based matchmaking system. In W. Lewis Johnson and Barbara Hayes-Roth, editors, *Proceedings of the 1st International Conference on Autonomous Agents*, pages 301–307, New York, February5–8 1997. ACM Press.

[FT91] D. Fudenberg and J. Tirole. *Game Theory*. The MIT Press, Cambridge, MA, 1991.

[GM98a] R. H. Guttman and P. Maes. Agent-mediated integrative negotiation for retail electronic commerce. In *Proc. of the Workshop on Agent Mediated Electronic Trading*, Minneapolis, Minnesota, May 1998.

[GM98b] R. H. Guttman and P. Maes. Cooperative vs. competitive multi-agent negotiations in retail electronic commerce. In *Proc. of the second Int. Conference on Cooperative Information Agents*, Paris, France, July 1998.

[GMM98] R. H. Guttman, A. Moukas, and P. Maes. Agent-mediated electronic commerce: A survey. *Knowledge Engineering Review*, June 1998.

[JBB+97] R. J. Bayardo Jr., B. Bohrer, R. S. Brice, A. Cichocki, J. Fowler, A. Helal, V. Kashyap, T. Ksiezyk, G. Martin, M. Nodine, M. Rashid, M. Rusinkiewicz, R. Shea, C. Unnikrishnan, A. Unruh, and D. Woelk. Infosleuth: Semantic integration of information in open and dynamic environments. In *Proc. of the ACM SIGMOD Conference on Management of Data*, pages 195–206, Tucson, Arizona, USA, 1997.

[JFN+96] N. R. Jennings, P. Faratin, T. J. Norman, P. O'Brien, M. E. Wiegand, C. Voudouris, J. L. Alty, T. Miah, and E.H. Mamdani. Adept: Managing business processes using intelligent agents. In *Proc. of the BCS Expert Systems 1996 Conference*, pages 5–23, 1996.

[KGM96] S. P. Ketchpel and H. Garcia-Molina. Making trust explicit in distributed commerce transactions. In *Proc. of the IEEE Int. Conference on Distributed Computing Systems*, 1996.

[KH96] D. Kuokka and L. Harada. Integrating information via matchmaking. *Journal of Intelligent Information Systems*, 6(2/3):261–279, 1996.

[NJFM96] T. J. Norman, N. R. Jennings, P. Faratin, and E.H. Mamdani. Designing and implementing a multi-agent architecture for business process management. In *Proc. of the third Int. Workshop on Agent Theories, Architectures, and Languages*, 1996.

[NM47] J. Von Neumann and O. Morgenstern. *Theory of games and economic behavior*. Princeton University Press, Princeton, New Jersey, second edition, 1947.

[Rai82] H. Raiffa. *The Art and Science of Negotiation*. Harvard University Press, Cambridge, Mass., 1982.

[RZ94] J. S. Rosenschein and G. Zlotkin. *Rules of Encounter: Designing Conventions for Automated Negotiation Among Computer*. The MIT Press, Cambridge, MA, 1994.

[San95] T. W. Sandholm. Limitations of the Vickrey auction in computational multiagent systems. In Victor Lesser, editor, *Proc. of the first Int. Conference on Multi-Agent Systems*. The MIT Press, Cambridge, MA, 1995.

[SFJ97] C. Sierra, P. Faratin, and N. R. Jennings. A service-oriented negotiation model between autonomous agents. *Lecture Notes in Artificial Intelligence*, 1237:17–35, 1997.

[SL95] T. W. Sandholm and V. R. Lesser. Coalition formation among bounded rational agents. In *Proc. of the 14th Int. Joint Conference on Artificial Intelligence*, pages 662–669, Montreal, Canada, 1995.

[TA98] T. Tesch and K. Aberer. Scheduling non-enforceable contracts among autonomous agents. In *Proc. of the third IFCIS Int. Conference on Cooperative Information Systems*, pages 86–94, New York City, USA, 1998.

[Tyg98] J. D. Tygar. Atomicity versus anonymity: Distributed transactions for electronic commerce. In *Proc. of the 24th Int. Conference on Very Large Databases*, pages 1–10, New York City, USA, 1998.

[Var95] Hal R. Varian. Economic mechanism design for computerized agents. In *First USENIX Workshop on Electronic Commerce*, pages 13–21, New York, July 1995. USENIX.

[Vic61] W. Vickrey. Counter specification, auctions, and competitive sealed tenders. *The Journal of Finance*, pages 9–37, 1961.

[ZR89] G. Zlotkin and J. S. Rosenschein. Negotiation and task sharing among autonomous agents in cooperative domains. In *Proc. of the 11th Int. Joint Conference on Artificial Intelligence*, pages 912–917, Detroit, Michigan, USA, August 1989.

[ZS96] D. Zeng and K. Sycara. How can an agent learn to negotiate? In *Proc. of the third Int. Workshop on Agent Theories, Architectures, and Languages*, 1996.