Advances in Complex Systems, Vol. 14, No. 2 (2011) 111–131

DOI: 10.1142/S0219525911003049

TRUST AS THE BASIS OF COALITION FORMATION IN ELECTRONIC MARKETPLACES

FRANK E. WALTER

Chair of Systems Design, ETH Zurich
Kreuzplatz 5, 8032 Zurich, Switzerland
fewalter@ethz.ch

Received 23 September 2010
Revised 6 March 2011

Despite the fact that social networks are ubiquitous on the Internet, only few websites exploit the potential of combining user communities and online marketplaces. Not many platforms allow users to engage in a phenomenon called “group buying” — buyers joining groups, or coalitions, to bundle their purchasing power towards sellers. We argue that this may be due to a lack of face-to-face interaction on the Internet; often, users do not know which other users to trust, which makes them suspicious of engaging in online business, in particular if many unknown other parties are involved. This situation, however, can be alleviated by leveraging the social networks of users: based on who a user knows and is connected to, a trust metric — for example, the TrustWebRank metric developed by us — can be computed to assess who else may be considered trustworthy to that user. In this paper, we build a simple agent-based model of coalition formation among agents in the setting of group buying in an electronic marketplace. In this model, agents use their trust relationships in order to determine who to form coalitions with. We show that this leads agents to experience high utility and that agents are able to learn who is trustworthy and who is not, even when they have no initial knowledge about the trustworthiness of other agents. This work may provide the foundation for a real-world application of an online coalition formation platform for e-commerce built on a social networking platform such as Facebook.

Keywords: Trust; social networks; coalition formation.

1. Introduction

One of the core features of the increasing ubiquity of the Internet is the fact that it provides the opportunity to form *spontaneous*, *location-independent communities* in different contexts. During the recent years, we have witnessed a boom of social networking websites such as MySpace, Facebook, and LinkedIn [5]. On such platforms, users have the possibility to communicate with other users, share pictures, videos, or links, and even organize themselves into groups. Such websites show increasing popularity, their user bases grow day by day. However, despite the ubiquity of such communities, only very few are used commercially, e.g. as marketplaces and for economic purposes. Of course, there are many websites which allow buying and selling products and services online: for example, Amazon or Ebay. However, only few websites exploit the potential of combining user communities and online marketplaces. This is surprising, in particular given the fact that there have been, across many cultures, time-proven traditions to bargain for prices and features of products and services in the off-line world. In many cultures, it is common for buyers and sellers to negotiate and bargain. Similarly, it is often common for buyers to join groups, or coalitions, to bundle their purchasing power towards sellers. So far, however, all this is not yet possible on the Internet. Usually, products or services are sold *at the price declared*, *with the features listed*, *and to one individual only*.

The proliferation of a marketplace principle originating in China seems to have the potential to change this. In Tuangou, several buyers form a coalition to negotiate with a seller for a special price or special features of a product or service. Literally, it translates to group buying, and it essentially is a process of coalition formation among agents which seek to maximize their utility by finding and cooperating with like-minded agents.

From an economic perspective, coalition formation may be a way to avoid having to compromise for a trade-off between *economies of scale* and *matching of preferences*. Economies of scale occur when mass production of standardized items reduces the costs of producers. Matching of preferences is achieved through the individual production of customized items for consumers. Clearly, there is a conflict between economies of scale and matching of preferences — this conflict leads to the fact that the diversification of products is mostly producer-driven in the sense that producers develop products that are then put on the market and evaluated by consumers in the process of market selection. This feedback from consumers again influences the diversification of products, i.e. it makes producers react by adjusting their supply to the demand of the market. Effectively, this leads to some products being eliminated from and others being introduced to the market.

Coalition formation among consumers provides another approach to influence the diversification of products: if there are enough customers that join, they can propose a product to the market that would suit their preferences. This would make the process of diversification of products more consumer-driven than it is at the moment. Of course, producers already rely on the efforts of market research to find out about the preferences of consumers. In addition to that, the information they can potentially receive through coalitions can serve them as another, inexpensive source of information about the market that they are trying to address. This is desirable both from the perspective of producers — as they better get to know the preferences of buyers and they have the chance to interact with consumers who might otherwise never buy from them — and from the perspective of consumers, as they can influence sellers, let them know their preferences in a better way, and if they are enough, bargain for discounts.

Thus, it appears that there are many arguments supporting that coalition formation and Tuangou-inspired marketplaces may be beneficial for both buyers and sellers. Why they have not yet become very popular among Internet users may be a question of *trust*: of course, it is straightforward to team up with communities of users online, but how is it possible to know whether these users are trustworthy, especially if money is involved? There is a history of attempts to scam users on the Internet, which is why it is understandable that users feel uncomfortable when dealing with strangers on the Internet [8]. By providing an infrastructure to reason about trust in such a system, we can design a trust-based mechanism of coalition formation. In our context [22, 23], trust is meant to be the expectancy of an agent to be able to rely on some other agent having similar preferences and, at the same time, to behave conforming to the rules for coalitions to successfully form. There has been a body of work on "trust webs" [1, 6, 17], and, in particular, on the issue whether and how trust between agents which do not know each other can be computed [13, 16]. In [23], we established the TrustWebRank metric for measuring trust in a social network and showed that it can be used as an alternative to established methods of generating recommendations in recommender systems. In this paper, one of our intentions is also to show that this metric — and possibly other related metrics that allow agents to infer trust in a complex system — can also be applied in a coalition formation context.

We introduce a very simple, basic model which is still far from reality in the sense of modeling, for example, a real electronic marketplace. However, throughout the paper, we will give our ideas of how to improve this model in the future and make it more realistic. The point we would like to make with the model is to show the basic applicability of metrics to compute trust in social networks in such a scenario of coalition formation in electronic marketplaces. We will show that *when users take trust relationships into account during coalition formation on social networks, they are able to team up with users that are trustworthy — even though they do not yet know them, i.e. even if they are far away in the social network.* This allows users to leverage the advantages of Tuangou (better match of preferences or better prices) without the disadvantages (uncertainty about who they are interacting with).

This paper was inspired by ideas in [21] and is an extension of a paper presented at the Fifth International Workshop on Emergent Intelligence on Networked Agents (WEIN 2010) at the Ninth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2010).

2. A Case Study: Store Mobs

In this section, we will shed light on a relatively recent and related phenomenon: *store mobs*. Unfortunately, this phenomenon has not yet been discussed in scientific literature, so all we can refer to are a number of articles in popular media. Store mobs, according to a platform that organizes store mobs, StoreMob.com, are "teams of people buying from a shop collectively to get a larger discount than if they shopped individually".[a] This sounds very similar to the idea of team buying; however, a store mob takes place in the physical world in a real store rather than on a website in the virtual world on the Internet:

> "Typically, the participants meet at the front of the targeted shop at a designated time. The group will then enter the store together, and the store mob organizer will talk to a staff member to explain the group would like a large discount (or other form of incentive such as free gifts) for any purchases they make then and there. The negotiations between the group and the store will then continue from there. Usually a larger discount can be achieved the greater the amount being spent, so big ticket purchases will tend to command better savings than smaller purchases." (StoreMob.com)

In other words, in a store mob, people organize themselves to show up at a particular store in groups. The larger the group, the better for the store mob. One of the participants of the store mob acts as an organizer who tries to negotiate some form of benefit with the store (in the form of discounts or other compensation) on behalf of the store mob participants. Usually, the organizer and the participants have already agreed on what terms they would like to achieve when they are at the store — for example, in a forum on the Internet. The leverage of the store mob during the negotiations are the surprise that they cause among the store and the number of participants. If the store agrees to a deal, the participants of the store mob buy the goods that they wanted to purchase and receive some form of benefit; if not, they all leave the store, possibly protesting in the process.

Of course, the participants have a benefit from such actions if the store enters a compromise: they get products and services at a better price (such as a 10% discount or "pay-for-one-take-two" offers) or with a free gift (such as a free USB mouse and keyboard in addition to a laptop). They do, however, also have a cost: they need to organize the store mob so that they can surprise the store. This requires effort and time, in particular when the size of the store mob gets large. However, there exist forums on the Internet which facilitate this process. Moreover, participants face the possibility that the store is not willing to offer a deal. In this case, all the participants of the store mob have to leave the store without exception, otherwise they lose their credibility.

The reasons why a store would negotiate with a store mob are twofold: on the one hand, it is natural to assume that a store would be willing to deal with potentially many customers who all came to buy something at the same time — they would lose business if they did not negotiate. Also, they may be able to make an offer which, although it does not meet the request of the store mob completely, still satisfies most of its participants. Such a compromise may make the customers happy, and happy customers return to the store. On the other hand, the store mob

[a] "How a Store Mob Works", retrieved from `http://storemob.com/how-a-store-mob-works/` on May 09, 2009.

also puts the store under a form of social pressure: if there is a group of, say, 500 customers[b] that all show up at the same location at the same time, this will probably frighten even very experienced store officials and create an unprecedented situation for them. Rather than facing the protests of so many people (possibly even publicly in the media or on the Internet) that would arise if they would briskly reject any negotiation, most of them will probably choose to concede and at least make some offer.

Nowadays, there exist platforms on the Internet which allow users to coordinate store mobs, e.g. StoreMob.com. These evolved from websites such as Mercata.com, LetsBuyIt.com, or MobShop.com that are all offline by now or have changed their business models. A few years ago, these websites served as platforms for more traditional team buying on the Internet in order to achieve volume discounts.

However, several articles[c] note that these websites never took off or went bankrupt because in fact, they did not offer consumers more ways of influencing the price or features of products and services: of course, such websites did bundle orders and give discounts based on how many customers were interested in a product or service. This apparently was part of a more sophisticated strategy: *price discrimination*, the sale of identical products or services from the same producer to different consumer segments at different prices. In this way, such team buying websites were merely used to attract customers who would otherwise not have bought particular goods or services.

Store mobs seem to have more potential than just being an instrument for price discrimination for sellers. This results from the fact that, in a store mob, buyers organize themselves (rather than relying on a seller-sponsored website). This gives buyers an edge during negotiation and makes store mobs so successful — in particular, in cultures where haggling has a century-old tradition, such as China.

3. Related Work

In the context of multi-agent systems, research on coalition formation among agents mostly focuses on coalitions as a means for distributed problem-solving rather than on the interactions between a seller and a group of buyers; the bulk of the literature mostly takes a game-theoretic approach and its perspective is on the "stability, fairness and payoff distribution" of coalitions of agents [20]. Coalition formation in an economic sense — i.e. group buying or buying clubs — has received by far less attention. In the following sections, we will give an overview of the literature on coalition formation among agents in electronic markets.

[b]See "Chinese consumer power: shop affronts" in the June 29th 2006 issue of The Economist (Volume 380, Number 8484, Page 6).
[c]"Does Group-Shopping Work?", Slate.com, July 26, 2000, and "When Teaming Up To Buy Pays Off", Money.com, June 1, 2000, both retrieved on May 09, 2009.

Overall, the core motivation seems to be the "*more buyers, lower cost*" principle which makes possible volume discounts the motivating incentive for both buyers and sellers being willing to deal with coalitions: greater volumes will lead to lower cost, and thus higher utility. The argument that coalition formation could also be used by buyers to achieve a critical mass when demanding custom features from sellers is only seldomly mentioned. [20] reports on "coalition formation as a means to formation of groups of customers coming together to procure products at a volume discount." They provide an incentive analysis for the formation of buying clubs, looking at the scenario both from the perspective of customers and suppliers. They establish utility functions for both buyers and sellers and, based on these, characterize the conditions which have to be met for both of them to engage in buying clubs. Additionally, [20] provides an overview of possible models which could be used to implement such coalition formation processes and they discuss their respective properties. [24] subsequently proposes a coalition formation scheme that essentially consists of a matching process between buyers and sellers based on the preferences of buyers and the prices offered by sellers. [24] formally analyzes their scheme and verify their findings with computer simulations. They claim that, because of the fact that their scheme is stable and efficient, it may be better suited for buying clubs online than existing systems. In particular, their approach is much more scalable than many of the existing, game-theoretic approaches — a necessary property for systems with potentially very many users on the Internet. [7] extends the concept of team buying to also account for "bundle search". Bundle search occurs when a set of products is to be purchased as a bundle — such as, for example, a computer together with a monitor. [7] proposes coalitions as a means of improving "bundle search" and design a distributed mechanism for doing so. Further, [19] discusses the impact of coalition formation on search cost. They argue that coalitions enable agents to share the cost of searching for a particular product: on the one hand, the search cost is split among many agents and on the other hand, the search horizon is widened the more agents are in the coalition. [19] argues that this could be another motivating incentive for agents to form coalitions.

A great share of the literature on coalition formation focuses on the design of effective mechanisms for coalition formation [3, 9–11, 14, 15, 18]. Some issues that have received attention in this context are:

- the fact that mechanisms have to be designed in such a way that they are efficient, yet agents have incentives to participate in coalitions [15];
- the fact that agents may not all have the same, common knowledge that often is assumed in coalition formation mechanisms and that this can drastically change their behavior and/or the properties of a coalition formation mechanism [11];
- the possibility that it is essential for agents to know who they form coalitions with and that agents use their trust relationships in order to determine who potential candidates to form a coalition with could be [3].

For a further overview of the literature please also [12].

4. A Simple Model

Suppose we have a system of agents which can purchase products and services in coalitions. Each agent i has one profile π_i of N_π different profiles that reflects its preferences.

Coalitions c are initiated by an initiator agent j; the initiator determines the profile π_c of the coalition, i.e. $\pi_c = \pi_j$. The intuition is that the initiator agent is an agent which — on behalf of all the other buyers in a coalition — acts as a mediator between the coalition and the seller, thus doing the bargaining for price and/or features. Let N_c denote the number of coalitions and $|c|$ the number of agents that have joined a particular coalition.

Agents i perceive a utility from joining a coalition, depending on whether the coalition has the same profile as they do or not. The objective of agents is to adjust their behavior of initiating or joining coalitions in such a way that their utility is maximized. To do so, they have to enter coalitions with like-profiled other agents as often as possible. We will start with a simple model, letting the utility function of an agent at time t be defined as

$$u_{ic}(t) = \begin{cases} 1 & \text{if } \pi_i = \pi_c \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

At a later time, it would be possible to (i) define profiles in such a way that it is possible to define a *distance* between any two profiles and to (ii) introduce a *threshold* on the size of coalitions, denoted N_{thr}, for them to be successful. These extensions would make the model more realistic.

If the distance between two profiles π_i and π_c is defined by a function $\Delta(\pi_i, \pi_c) \in [0, 1]$, where a value of zero signifies identical profiles and a value of one signifies opposite profiles, the utility function could be defined as

$$u_{ic}(t) = 1 - \Delta(\pi_i, \pi_c). \tag{2}$$

As there can be more similar and less similar profiles, this may potentially lead to heterogeneous coalitions if there is a lack of options for agents. Furthermore, let us consider another constraint: if not at least N_{thr} agents have joined a particular coalition, the coalition fails and none of the agents perceives a utility. Combining all these extensions, the utility function could then be defined as

$$u_{ic}(t) = \begin{cases} 1 - \Delta(\pi_i, \pi_c) & \text{if } |c| \geq N_{\text{thr}} \\ 0 & \text{otherwise.} \end{cases} \tag{3}$$

This would mimic the situation that, in reality, there needs to be a certain number of agents in a coalition for the buyers' bargaining to pay off and for the sellers to be willing to bargain at all. If $|c| < N_{\text{thr}}$, agents will only join a coalition if they are confident that the initiator will eventually be able to motivate at least N_{thr} agents. For the moment, we will, however, stay with the case of Eq. (1).

In order to be able to compare the utilities across all agents at different time steps t, we define the performance of the system Φ as

$$\Phi(t) = \frac{1}{n} \sum_i u_{ic}(t), \tag{4}$$

where n is the number of agents.

Obviously, if agents do not have any knowledge about other agents, they do not know which coalitions to join. In this case, all they can do is to choose randomly among all coalitions. If there are N_π profiles in the system, the probability of joining a coalition that is of the same profile as of the joining agent is $\frac{1}{N_\pi}$. In all other cases, the joining agent is joining a coalition of a different profile, leading to it not experiencing the optimal utility.

Of course, it would be possible for an agent i to remember, for each other agent j, whether this agent had the same profile or not. However, we assume that — to model a real system — this is not possible since there are too many agents to remember the profile of each one and since profiles of agents may change over time (and in this way, information may become outdated).

Furthermore, remembering the profiles of all other agents also does not solve the problem: if the system is large, most agents will not be able to interact with all possible partners, and thus only know the profiles of a small number of other agents. From that knowledge, they cannot deduce the profiles of the other agents that they do not yet know. In other words, even though an agent knows the profile of some coalitions, it does not know the profile of many coalitions and thus it is not able to select an optimal one.

All this changes if there is a way for agents to select the coalition they join based on the level of trust that they have in the initiator of the coalition. If this is the case, agents can distinguish between initiators and thus between coalitions which are like-minded and those that are not. Even if, initially, the level of trust between agents is not known, agents can dynamically learn what trust they should have in neighbors. In other words, subsequent to a phase of learning, they will be able to tell which agents are trustworthy and which ones are not. Since trust in the system evolves dynamically, the system is robust also with respect to changing profiles among agents.

Let the system of agents embedded in a social network be defined by a graph. The graph is associated to an adjacency matrix A in which $A_{ij} = 1$ if two agents i and j are neighbors and $A_{ij} = 0$ otherwise. Each agent i keeps track of its trust relationships to its neighbors j. These are reflected in the matrix of *direct trust* T with entries $T_{ij} \in [0, 1]$. A value of $T_{ij} = 1$ signifies full trust of agent i in agent j, a value of $T_{ij} = 0$ signifies no trust of agent i in agent j. It may happen that an agent has a neighbor that it does not trust, i.e. $A_{ij} = 1$ and $T_{ij} = 0$. However, $T_{ij} \neq 0$ only if $A_{ij} = 1$, i.e. the matrix T expresses trust only between neighbors.

The number of neighbors of an agent is referred to as its degree. The average number of neighbors across all agents is referred to as the average degree d of agents.

In order to compute the *indirect trust* $\tilde{T}$ between agents which are not neighbors, we can apply the TrustWebRank metric. In analogy to PageRank, this metric is inspired by the concept of feedback centrality which assigns a centrality score to the nodes of a network based on the scores of the node's neighbors: in a trust context, the higher (or lower) the trustworthiness of a node's neighbors, the higher (or lower) this node's own trustworthiness [23]. Unlike PageRank, TrustWebRank is personalized, i.e. it determines "how trustworthy agent i is from the perspective of agent j" [23]; the trustworthiness of an agent is not a *global* value:

$$\tilde{T}_{ij} = S_{ij} + \beta \sum_{k \in N_i} S_{ik} \tilde{T}_{kj} \quad \forall\, i, j, \tag{5}$$

where N_i is the set of neighbors of agent i, $\beta \in [0, 1)$ is a dampening factor (we choose $\beta = 0.8$, see [23]), and S_{ij} is the normalised representation of $T_{ij}, S_{ij} = \frac{T_{ij}}{\sum_{k \in N_i} T_{ik}}$.

This implies that we are able to determine the direct trust between agents that are neighbors and the indirect trust between those which are not neighbors. This knowledge could be used by agents in order to determine which coalition they should join: since the profile of a coalition is determined by the initiator, agents should choose coalitions which have a trustworthy initiator. If there are no such coalitions, they should initiate their own coalition.

This suggests that it would be possible that any agent simply forms their own coalition, regardless of what other agents do and which coalitions already exist. In real life, this would not happen because an initiator usually has more duty (negotiation and bargaining) and responsibility (acting as the interface between buyers and sellers), and thus there also is a cost to being an initiator. Additionally, coalitions may fail if there are not enough members, leading to a utility of 0. We will incorporate these aspects in a future version of the model.

Thus, we have two approaches: agents choosing their coalitions at random, or choosing their coalitions based on the level of trust to initiators.

4.1. *Random choice of coalitions*

When agents choose their coalitions at random, they follow the decision tree in Fig. 1 to determine their behavior at each time t.

According to this structure, at each time t, agents proceed in several rounds r until each of the agents has either initiated a coalition or chosen one to join.

As long as they have not yet initiated or joined a coalition, during each round, agents decide whether they should immediately act or wait until the next round. This decision is stochastic, where the probability to wait follows an exponential

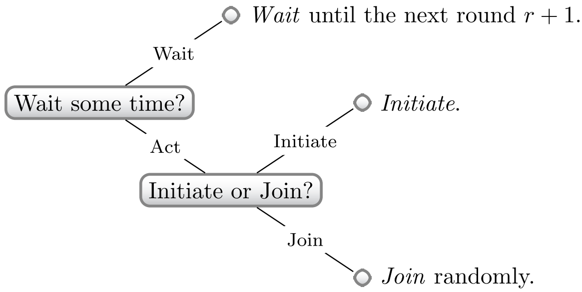

Fig. 1. Decision tree: random choice of coalitions.

decay function with

$$P(\text{wait}) = \exp(-(1-\delta)r), \quad P(\text{act}) = 1 - P(\text{wait}), \tag{6}$$

where $r = 1, 2, 3, \ldots$ is the current round and $\delta \in [0, 1]$ is a parameter controlling the exponential decay function. We use $\delta = 0.8$. Eventually, as r is increasing, $P(\text{wait})$ will get close to zero and $P(\text{act})$ close to one. Thus, at some point, it is very likely that all agents will have acted. The purpose of this step is that agents do not all choose their coalitions at once and that they do so in random order.

Once an agent has decided not to wait any more, but to act, it makes a choice to either initiate a coalition or to join any of the coalitions at random. This choice is stochastic, with the probabilities of initiating or joining being

$$P(\text{initiate}) = 1 - \varepsilon, \quad P(\text{join}) = \varepsilon, \tag{7}$$

where $\varepsilon \in [0, 1]$ is a parameter. Of course, if there is no coalition yet, an agent has no choice but to initiate one, as there is none to join.

Of course, a random choice of coalitions is not a smart strategy for agents to follow. However, if an agent has no information about other agents, this is the only strategy that it can follow. The random choice of coalitions is a benchmark for evaluating other, more sophisticated strategies.

4.2. *Trust-based choice of coalitions*

When agents choose their coalitions based on the trust that they have in them, they follow the decision tree in Fig. 2 to determine their behavior at each time t.

Again, agents proceed in rounds r and potentially wait some time until they act — at some later time, there may be more coalitions with agents that are trustworthy as initiators. Once an agent has decided not to wait any more, but to act, it tries to find out whether there are any trustworthy coalitions or not. If there are some, it will select one of them at random. We could extend this model at a later time by changing this behavior and introducing some form of preferential selection in the sense that agents select among trustworthy coalitions with probabilities

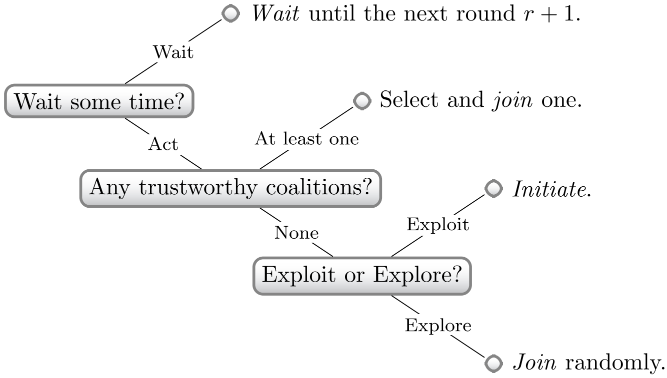

Fig. 2. Decision tree: trust-based choice of coalitions.

weighted according to:

- the trust that they have in the initiator;
- the trust that they have into all the other participants of a coalition;
- the number of neighbors among the other participants of a coalition;
- the number of participants of a coalition in case of a threshold of required agents for a coalition not to fail.

If an agent does not see any trustworthy coalitions, this can have two reasons: either (i) there really are none, or (ii) the agent has not yet developed trust to the initiator. Agents cannot distinguish between these two cases — this represents an exploration–exploitation dilemma. If agents would try two, three coalitions and then stick to the one that they did best with or initiate their own coalition all the time (this would be exploitation only), they risk not seeing better options, potentially leading them to not optimal performance. Equally, if agents would continuously try out different coalitions (this would be exploration only), they would probably do well sometimes and find good coalitions, but also do badly at other times, leading them to not optimal performance overall, too. The optimum lies in a trade-off between exploration and exploitation.

In this context, exploration would be to choose a coalition at random and exploitation would be to initiate a coalition by themselves — which is known to lead to good performance. Note that this choice is — by the decision tree that agents follow — only made when there are no other trustworthy coalitions which would otherwise be the choice for exploitation.

The current version of the model can, in this regards, be improved. If, for example, there were a cost for agents to initiate a coalition or the utility of a coalition grew proportionally to its size (the sequence of arguments is the following: the more agents, the better the bargain, the higher the utility), initiating a coalition would not necessarily be a better choice than choosing at random. We plan to incorporate these aspects into a future version of the model.

It is possible to achieve a trade-off between exploration and exploitation by making a stochastic choice in which

$$P(\text{initiate}) = 1 - \varepsilon, \quad P(\text{choose at random}) = \varepsilon, \tag{8}$$

where $\varepsilon \in [0, 1]$ fulfils a similar function as in Eq. (7). In this way, agents explore their options, but also exploit ones which they know to lead to good performance. Once an agent has, during round r at time t, chosen to initiate or join a coalition, it waits until all other agents have made their choice.

5. Trust Dynamics

Assuming matrices T and $\tilde{T}$ reflecting the direct and indirect trust (and in that sense, the like-mindedness of agents in terms of profiles) between agents to be given, it is trivial to see that this approach will lead to agents that have trust in each other to form coalitions. This will be illustrated further on when we analytically approximate the behavior of the system and when we illustrate it by means of computer simulations. However, it is not trivial to see that by introducing trust dynamics, it is possible for agents to learn which other agents to trust as initiators of coalitions.

Suppose that an agent has participated in a coalition and is able to communicate with its co-participants. Based on its own utility and the utility that its co-participants report to perceive, the agent is able to update its trust to its neighbors (not everyone in the population) as follows:

$$T_{ij}(t+1) = \begin{cases} T_{ij}(t) & \text{if } u_{ic}(t) = u_{jc}(t) = 0; \\ \gamma T_{ij}(t) + (1-\gamma)(1 - |u_{ic}(t) - u_{jc}(t)|) & \text{otherwise,} \end{cases} \tag{9}$$

where γ is a parameter controlling the memory size of agents. Note that these trust dynamics are, in a sense, very similar to the ones in [23]; following this reference, we choose $\gamma = 0.8$. The idea of this approach is that agents develop local knowledge about which of their neighbors are trustworthy or not. By applying the TrustWeb-Rank metric and computing $\tilde{T}$ from T, they can compute how trust propagates in the social network and apply this knowledge globally in order to find trustworthy initiators.

At the moment, in this very simple, basic model, trust between two agents is thus based on how similar their profiles are. In future versions of the model, more aspects than just profile similarity may be incorporated — for example, what an agent thinks of another agent's bargaining skills, to what extent an agent contributed to the success of a coalition, whether an agent paid or delivered on time, and many more.

6. Analysis of the Model

In this section, we will derive a characterization of the behavior of our model by a mean-field approximation. On the one hand, we will compute analytical approximations for the performance of agents when choosing coalitions at random, without being able to take trust relationships in their social network into account. On the other hand, we will show that, by first, using trust to identify coalitions to join and by second, exchanging feedback about the utility perceived in particular coalitions with neighbors, performance of agents approaches the optimum.

Suppose that agents choose their coalitions at random, following the procedure described in Sec. 4.1 when they do not yet have any information about who is trustworthy and who is not. When following this procedure, each agent initiates a coalition with probability $1-\varepsilon$ and it joins a random coalition with probability ε — during which round an agent makes it decision does not affect these probabilities. When an agent initiates a coalition, its own profile is guaranteed to match the profile of the coalition, hence its utility will be one. When an agent joins a random coalition, only $\frac{1}{N_\pi}$ of the coalitions will have the same profile, i.e. its utility will be zero except in these (few) cases. Thus, the expected utility of an agent i at time t, denoted by $u_i := E(u_{ic}(t))$, can be expressed as

$$u_i = (1-\varepsilon)1 + \varepsilon\frac{1}{N_\pi}. \tag{10}$$

Since u_i is the same for all agents, and the performance of the system $\Phi(t)$ is defined as the average of all utilities $u_{ic}(t)$ of all agents at time t, the expected performance of the system $\Phi := E(\Phi(t)) = u_i = (1-\varepsilon)1 + \varepsilon\frac{1}{N_\pi}$. For example, given $\varepsilon = 0.8$ and $N_\pi = 2$, $\Phi = 0.6$. Note that, if agents do not yet know who is trustworthy, they follow this behavior — hence, this can also be seen as an approximation of the baseline performance of the trust-based approach!

Now suppose that agents choose their coalitions based on the trust they have in them (i.e., in initiators, respectively), following the procedure described in Sec. 4.2. It is clear from the decision tree that, once agents know who is trustworthy, they join coalitions with trustworthy initiators if they can do so. Hence, what we need to show is that agents learn, through the trust dynamics, which of their neighbors are trustworthy and which ones are not. We already showed in [23] that, given a direct trust matrix, the TrustWebRank metric allows agents to reliably compute an indirect trust matrix.

Thus, we need to show that the trust dynamics of Eq. (9) converge to a state in which two agents i, j of the same profile trust each other and agents of different profiles do not trust each other; this leads to the following three cases:

Case 1: If two agents i, j join the same coalition c and both experience a positive utility ($u_{ic}(t) = u_{jc}(t) = 1$, i.e. both have the same profile), the trust between them evolves as follows:

$$T_{ij}(t+1) = \gamma T_{ij}(t) + (1-\gamma)(1 - |u_{ic}(t) - u_{jc}(t)|). \tag{11}$$

We need to show that this sequence converges to 1, i.e. that given $u_{ic}(t) = u_{jc}(t) = 1$,

$$\lim_{t\to\infty}\left(\gamma T_{ij}(t) + (1-\gamma)(1 - \underbrace{|u_{ic}(t) - u_{jc}(t)|}_{1-1=0})\right) = 1. \tag{12}$$

Let $a_t := T_{ij}(t) - 1$, so $T_{ij}(t) = a_t + 1$ and $T_{ij}(t+1) = a_{t+1} + 1$. We can substitute this into Eq. (11) and derive a recursion for a_t:

$$a_{t+1} + 1 = \gamma(a_t + 1) + (1-\gamma)1 = \gamma a_t + 1, \quad \text{i.e.} \tag{13}$$

$$a_{t+1} = \gamma a_t. \tag{14}$$

Therefore, $a_t = \gamma a_{t-1}$. Substituting this in Eq. (14), we get: $a_{t+1} = \gamma^2 a_{t-1} = \gamma^3 a_{t-2} = \cdots$ and so we derive an explicit expression for a_t:

$$a_{t+1} = \gamma^{t+1} a_0, \quad \text{or, for } a_t: \ \ a_t = \gamma^t a_0. \tag{15}$$

Resubstituting $a_t = T_{ij}(t) - 1$, we get

$$T_{ij}(t) = \gamma^t (T_{ij}(0) - 1) + 1. \tag{16}$$

From this, since $\gamma \in (0,1)$ and therefore $\lim_{t\to\infty} \gamma^t = 0$, it is clear that for $t \to \infty$, $\lim_{t\to\infty} T_{ij}(t) = 0 + 1 = 1$, independent of the value of $T_{ij}(0) \in [0,1]$. This shows that the trust between two agents i, j that have the same profile converges to one.

Case 2: If two agents i, j join the same coalition and both experience a negative utility (i.e. both have a different profile as the initiator, but not necessarily the same one), the trust between them does not change. This is appropriate, as they may, but do not have to have different profiles. None of the agents is able to determine whether it should increase or decrease its trust to the neighbor, hence they do not do anything.

Case 3: If two agents i, j join the same coalition and one of them experiences a positive, and the other one a negative utility ($u_{ic}(t) \neq u_{jc}(t)$, i.e. both have different profiles), the trust between them evolves as follows:

$$T_{ij}(t+1) = \gamma T_{ij}(t) + (1-\gamma)(1 - |u_{ic}(t) - u_{jc}(t)|). \tag{17}$$

We need to show that this sequence converges to 0, i.e. that given $u_{ic}(t) \neq u_{jc}(t)$,

$$\lim_{t\to\infty}\left(\gamma T_{ij}(t) + (1-\gamma)(1 - \underbrace{|u_{ic}(t) - u_{jc}(t)|}_{|0-1|=1 \text{ or } |1-0|=1})\right) = 0. \tag{18}$$

Let $b_t := T_{ij}(t)$. Then Eq. (17) gives us a recursion for b_t:

$$b_{t+1} = \gamma b_t. \tag{19}$$

Therefore, $b_t = \gamma b_{t-1}$. Substituting this in Eq. (19), we get: $b_{t+1} = \gamma^2 b_{t-1} = \gamma^3 b_{t-2} = \ldots$, and so we derive an explicit expression for b_t:

$$b_{t+1} = \gamma^{t+1} b_0, \quad \text{i.e. for } b_t: \;\; b_t = \gamma^t b_0. \tag{20}$$

Resubstituting $b_t = T_{ij}(t)$, we get

$$T_{ij}(t) = \gamma^t (T_{ij}(0)). \tag{21}$$

From this, since $\gamma \in (0,1)$ and therefore $\lim_{t\to\infty} \gamma^t = 0$, it is clear that for $t \to \infty$, $\lim_{t\to\infty} T_{ij}(t) = 0$, independent of the value of $T_{ij}(0) \in [0,1]$. This shows that the trust between two agents i, j that have a different profile converges to zero.

Thus, we have shown that the trust dynamics lead to trust between agents of the same profile and no trust between agents of different profiles.

Now suppose that agents initially do not have any knowledge about who to trust and who not to trust, i.e. $T_{ij}(0) = 0 \; \forall\, i, j$. Then, during the initial stages, agents will not find any trustworthy coalitions (as they do not trust anyone), but rather either initiate a coalition or join a coalition at random, just as if they did not base their decisions on trust at all. During this phase, agents do, however, gradually learn which of the neighbors are trustworthy. At later stages, agents will then know who is trustworthy and almost only join coalitions with trustworthy initiators, leading to an expected utility of $u_i \approx 1$ for almost all agents, which again implies an expected performance of the system $\Phi \approx 1$.

7. Simulations

The simulations we carried out were done on an agent population of 500 agents. The agents are connected in a random graph [2, 4]. Initially, $T_{ij} = 0 \; \forall\, i, j$, i.e. the agents have to learn who to trust; we chose $\gamma = 0.8$ for the trust dynamics (based on [23]). We varied the average degree d of each agent, as well as the number of profiles N_π in the system. The figures illustrate the system behavior over 300 steps; all results were averaged over 100 runs.

Figure 3 illustrates the proportion of agents that choose coalitions at random, initiate coalitions, and that choose coalitions based on trust over time. In one plot, the average degree of agents is fixed, $d = 30$, but the number of profiles is variable, $N_\pi \in \{2, 4, 6\}$; in the other plot, the average degree of agents is variable, $d \in \{10, 20, 30\}$, but the number of profiles is fixed $N_\pi = 2$. We can see that in both cases, the proportion of agents that choose coalitions based on trust increases rapidly over time; at the same time, both the proportion of agents that initiates coalitions and the proportion of agents that chooses coalitions at random decreases. For an increasing number of profiles (with a fixed average degree of agents), this process takes place more slowly. In other words, the more profiles, the more difficult it is for agents to learn who is trustworthy, and the slower the learning process takes place. For an increasing average degree of agents (with a fixed number of profiles), this process takes place more quickly. In other words, the the greater the average

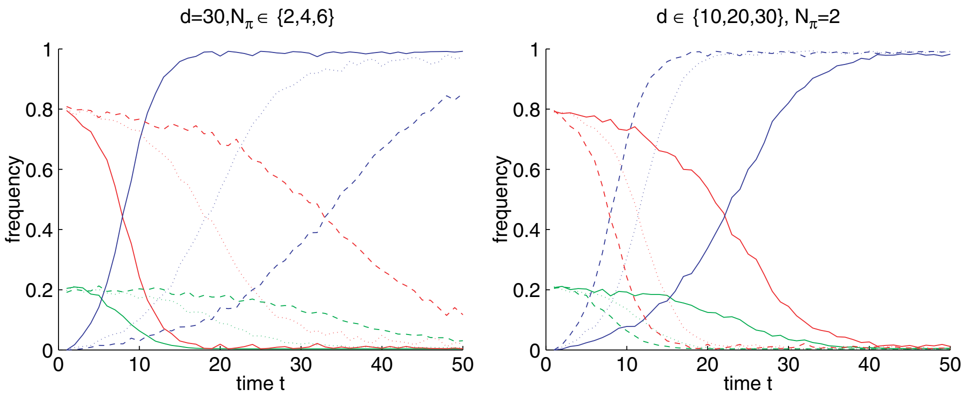

Fig. 3. (Color online) Coalition behavior over time, $t = 1, \ldots, 50$, for a fixed average degree of agents but variable number of profiles (left, solid lines $N_\pi = 2$, dotted lines $N_\pi = 4$, dashed lines $N_\pi = 6$) and for a variable average degree of agents, but a fixed number of profiles (right, solid lines $d = 10$, dotted lines $d = 20$, dashed lines $d = 30$). The plots show the proportion of agents that choose coalitions at random (red, $t = 0$-intersect at 0.8 and curves decreasing), initiate coalitions (green, $t = 0$-intersect at 0.2), and that choose coalitions based on trust (blue, $t = 0$-intersect at 0 and curves increasing).

degree of agents, the more likely it is that agents have a trustworthy neighbor of the same profile among their neighbors, and thus the faster they find these neighbors. Thus, we can see that, initially, agents join a lot of coalitions at random (or they initiate them) because they are not able to distinguish between "good" ones and "bad" ones. After some time, however, agents have learnt who to trust and they convert that knowledge into decisions of who to join; subsequently, they join coalitions based on trust.

Figure 4 illustrates the process of coalition formation at different time steps t. The nodes correspond to agents, the links indicate which initiator a particular agent joins in a coalition. In order for the process of coalition formation to be visualisable, we consider only 50 agents with two profiles, represented in blue and red color, respectively.

Figure 4(a) illustrates the case in which agents have not yet built trust relationships and thus choose randomly. It can be seen that several coalitions of various sizes emerge. Since agents choose randomly (they do not yet know who they can trust), they also do not form coalitions with like-profiled other agents. This also implies that the performance of the system will not be anywhere close to one as there will be many agents who have, from their perspective, joined the wrong coalition.

Next, Fig. 4(b) illustrates the case in which the agents are still learning, but have built trust to some agents. Agents mostly form coalitions with initiators that they know to be trustworthy, but in some cases, they also still choose randomly.

Finally, Fig. 4(c) illustrates the case in which agents have built trust relationships. Agents only form coalitions with initiators that they know to be trustworthy. As there are only two profiles, two giant coalitions, one for each profile, emerge.

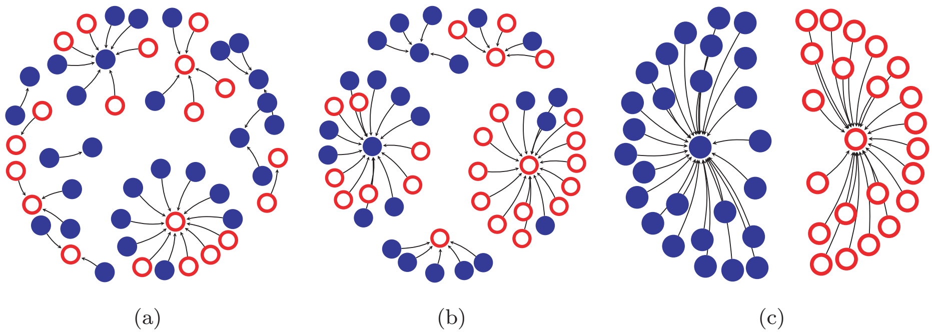

Fig. 4. (Color online) Different stages of the coalition formation process following the trust-based approach at $t = t_{\text{beginning}}$ (a), $t = t_{\text{intermediate}}$ (b), and $t = t_{\text{equilibrium}}$ (c).

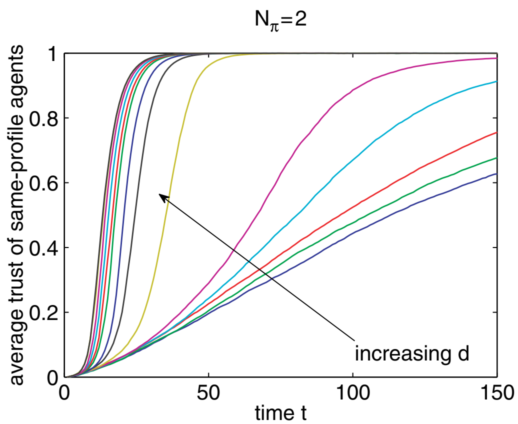

Fig. 5. Trust between agents of the same profile over time, for a variable average degree of agents, but a fixed number of profiles.

Note, however, that this need not necessarily be the case and that it would also be possible that several smaller coalitions emerge, but they all would be dominated by agents of a particular profile. This also implies that the performance of the system will be one as all agents have, from their perspective, joined the right coalition.

Figure 5 illustrates the trust between agents of the same profile over time: the average degree of agents is variable, $d \in [0, 50]$, and the number of profiles is fixed, $N_\pi = 2$. Also here the average trust between agents of the same profile converges to 1 for almost all d. However, there are some d for which this only takes place very slowly. Given a low average degree of agents, many agents are not connected to any other agents of their profile. In that way, they cannot build trust to anyone, and thus they are forced to choose their coalitions at random. Above a particular threshold ($d \approx 5$ for $N_\pi = 2$), agents have enough neighbors such that at least one of them has the same profile. In another case (not pictured), the average degree of agents is fixed, $d = 30$, and the number of profiles is variable, $N_\pi \in \{2, 3, 4, 5, 6, 7, 8, 9, 10\}$. The average trust between agents of the same profile converges to 1 for all N_π. For smaller N_π, this process takes much faster than for larger N_π. However, for all numbers of profiles, agents develop trust to other agents that have the same profile. Furthermore (this cannot be seen from the figure), agents of opposite profiles do not develop trust between each other.

Figure 6 illustrates the trust between agents of the same profile as a function of the number of profiles and average degree of agents at $t = 25$ and $t = 150$. Trust develops slowly compared to the model discussed in [23] because there are much less interactions with neighbors at each step t which can be used for feedback for the trust dynamics. Trust develops most quickly the smaller the number of profiles and the larger the average degree of agents. There is a threshold for d which depends on N_π and above which agents develop trust to other agents of the same profile.

Finally, Fig. 7 then shifts our focus to performance and illustrates the performance over time: the average degree of agents is fixed, $d = 30$, and the number of profiles is variable, $N_\pi \in \{2, 3, 4, 5, 6, 7, 8, 9, 10\}$. The performance converges to 1 for all N_π. For smaller N_π, this process takes much faster than for larger N_π. In another case (not pictured), the average degree of agents is variable, $d \in [0, 50]$, and the number of profiles is fixed, $N_\pi = 2$. Also here the performance converges to 1 for almost all d. However, there are some d for which this is not the case. Given a low average degree of agents, many agents are not connected to any other agents of their profile. In that way, they cannot build trust to anyone, and thus they are forced to choose their coalitions at random, which leads the system performance not to converge to 1. Above a particular threshold ($d \approx 5$ for $N_\pi = 2$), agents have enough neighbors such that at least one of them has the same profile, and as soon as this threshold is reached, the performance of the system converges to 1. The similarity to Fig. 5 is due to the fact that agents who have developed trust to

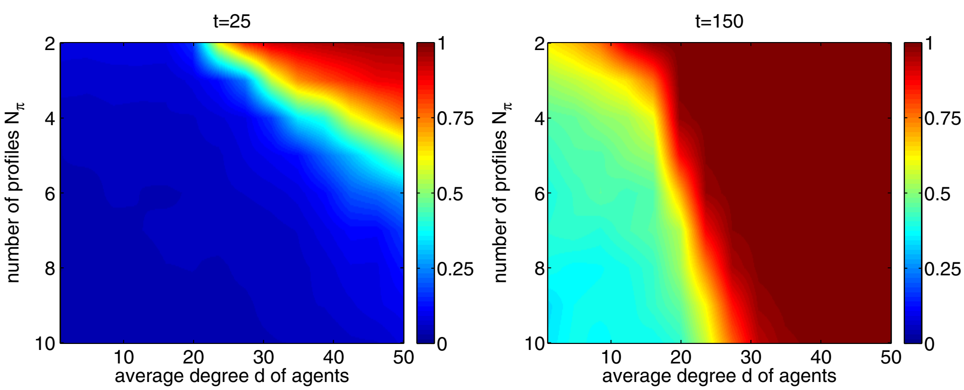

Fig. 6. Trust between agents of the same profile as a function of the number of profiles and the average degree of agents at various times t.

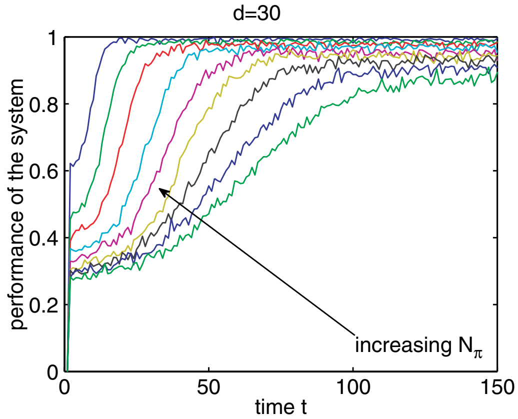

Fig. 7. Performance over time, for a fixed average degree of agents but variable number of profiles.

other agents of the same profile choose to join coalitions of initiators with the same profile; thus, these agents perceive high utility which leads to high performance.

8. Discussion and Conclusion

We applied the TrustWebRank metric to a scenario of coalition formation and built a simple model in which agents use their trust relationships in order to determine who to form coalitions with. We showed that — also in this scenario — this leads agents to good performance and that agents are even able to learn who is trustworthy and who is not given that they do not have any initial knowledge about the trustworthiness of other agents.

This shows, with a simple model, that it may, in fact, be possible to solve the problem of assessing the trustworthiness of strangers that one would eventually run into when combining user communities and on-line marketplaces. While such models of "trust webs" [1, 6, 17] have mostly been applied in different scenarios, such as recommender systems, their possible areas of application are very vast.

Some extensions to this model could involve changing the utility function:

Thresholds for coalition success and failure. It is unrealistic to assume that even small coalitions of buyers are able to negotiate with sellers. Therefore, there could be a threshold on the number of agents in a coalition which determines the success of a coalition. If agents join a failing coalition, they do not have any utility from that coalition.

Continuous profiles. At the moment, profiles are discrete — i.e., if two agents are of the same profile, their preferences match in every aspect. We could introduce continuous profiles and thus the possibility that two profiles have partial overlap. Depending on how much two profiles overlap, the distance between them varies.

Nonetheless, we showed that when users take trust relationships into account during coalition formation on social networks, they are able to team up with users that are trustworthy — even though they do not yet know them, i.e. even if they are far away in the social network. This allows users to leverage the advantages of group buying (better match of preferences or better prices) without the disadvantages (uncertainty about who they are interacting with).

References

[1] Abdul-Rahman, A. and Hailes, S., Supporting trust in virtual communities, in *Proceedings of the 33th Annual Hawaii International Conference on System Sciences* (IEEE Press, 2000), p. 6007.
[2] Bollobás, B., *Random Graphs* (Academic Press, 1985).
[3] Breban, S. and Vassileva, J., Using inter-agent trust relationships for efficient coalition formation, in *Proceedings of the 15th Conference of the Canadian Society for Computational Studies of Intelligence — Advances in Artificial Intelligence (AI 2002)*, eds. Cohen, R. and Spencer, B., *Lecture Notes in Artificial Intelligence*, Vol. 2338 (Springer, 2002), pp. 221–236.
[4] Erdős, P. and Rényi, A., On random graphs, *Publicationes Mathematicae Debrecen* **6** (1959) 290–291.
[5] Golder, S., Wilkinson, D. M. and Huberman, B. A., Rhythms of social interaction: Messaging within a massive online network, *3rd International Conference on Communities and Technologies (CT2007)*, East Lansing, MI, June 28–30 (2007).
[6] Grandison, T. and Sloman, M., A survey of trust in internet applications, *IEEE Communications Surveys and Tutorials* **3** (2000).
[7] He, L. and Ioerger, T. R., Combining bundle search with buyer coalition formation in electronic markets: A distributed approach through negotiation, in *Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'04)* (ACM Press, 2004), pp. 1440–1441.
[8] Hoffman, D. L., Novak, T. P. and Peralta, M., Building consumer trust online, *Commun. ACM* **42** (1999) 80–85.
[9] Indrawan, M., Kijthaweesinpoon, T., Srinivasan, B. and Sajeev, A. S. M., Coalition formation protocol for E-commerce, in *Proceedings of Intelligent Sensing and Information Processing* (2004), pp. 403–408.
[10] Ito, T., Ochi, H. and Shintani, T., A group buy protocol based on coalition formation for agent-mediated E-commerce, *Int. J. Comput. Inf. Sci.* **3** (2002) 11–20.
[11] Ketchpel, S., Forming coalitions in the face of uncertain rewards, in *Proceedings of the 12th National Conference on Artificial Intelligence* (AAAI Press, 1994), pp. 414–419.
[12] Klusch, M. and Gerber, A., Dynamic coalition formation among rational agents, *IEEE Intel. Syst.* **17** (2002) 42–47.
[13] Kuter, U. and Golbeck, J., Using probabilistic confidence models for trust inference in Web-based social networks, *ACM Trans. Internet Technol.* **10** (2010) 8:1–8:23.
[14] Lerman, K. and Onn Shehory, Coalition formation for large-scale electronic markets, in *Proceedings of the 4th International Conference on Multiagent Systems (ICMAS 2000)* (IEEE Computer Society, 2000), pp. 167–174.
[15] Li, C., Rajan, U., Chawla, S., and Sycara, K., Mechanisms for coalition formation and cost sharing in an electronic marketplace, in *Proceedings of the 5th International Conference on Electronic Commerce (ICEC '03)* (ACM Press, 2003), pp. 68–77.
[16] Liu, G., Wang, Y. and Orgun, M., Optimal social trust path selection in complex social networks, in *AAAI Conference on Artificial Intelligence* (2010).
[17] Sabater, J. and Sierra, C., Review on computational trust and reputation models, *Artif. Intel. Rev.* **24** (2005) 33–60.

[18] Sarne, D. and Kraus, S., The search for coalition formation in costly environments, in *Proceedings of the 7th International Workshop on Cooperative Information Agents (CIA 2003)*, eds. Klusch, M., Ossowski, S., Omicini, A., and Laamanen, H., *Lecture Notes in Computer Science*, Vol. 2782 (Springer, 2003), pp. 117–136.

[19] Sarne, D. and Kraus, S., Buyer's coalition for optimal search, in *Proceedings of the 4th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'05)* (ACM Press, 2005), pp. 1225–1226.

[20] Tsvetovat, M., Sycara, K., Chen, Y. and Ying, J., Customer coalitions in the electronic marketplace, in *Proceedings of the 7th Conference on Artificial Intelligence (AAAI-00) and of the 12th Conference on Innovative Applications of Artificial Intelligence (IAAI-00)* (AAAI Press, 2000), pp. 1133–1134.

[21] Walter, F. E., Battiston, S. and Schweitzer, F., Emergence and evolution of coalitions in buyer-seller networks, in *Proceedings of the Workshop on Emergent Intelligence of Networked Agents at the Fifth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'06)* (2006), published in Studies in Computational Intelligence (56), pp. 245–258, Springer, 2007.

[22] Walter, F. E., Battiston, S. and Schweitzer, F., A model of a trust-based recommendation system on a social network, *J. Autonomous Agents and Multi-Agent Systems* **16** (2008) 57–74.

[23] Walter, F. E., Battiston, S. and Schweitzer, F., Personalised and dynamic trust in social networks, in *RecSys '09: Proceedings of the third ACM conference on Recommender systems* (ACM Press, 2009), pp. 197–204.

[24] Yamamoto, J. and Sycara, K., A stable and efficient buyer coalition formation scheme for e-marketplaces, in *Proceedings of the 5th International Conference on Autonomous Agents*, eds. Müller, J. P., Andre, E., Sen, S., and Frasson, C. (ACM Press, 2001), pp. 576–583.