Game-Theoretic Recommendations: Some Progress in an Uphill Battle

Moshe Tennenholtz
Technion – Israel Institute of Technology
Haifa 32000, Israel

1. INTRODUCTION

Game theory has become the central language for the analysis of multi-agent systems. Moreover, the central game-theoretic solution concept, the Nash equilibrium, has become a standard tool for that analysis. A game is a general way for representation of interactions among agents: each agent has strategies he can choose from, and each tuple of strategies, one for each agent, determines a payoff for each of the agents. A Nash equilibrium is a strategy profile, such that unilateral deviations from it are not beneficial. However, this concept does not provide a solution to what we believe to be the major challenges of game theory and the theory of multi-agent systems:

1. Given a game, how *should* the agent choose his action?
2. Given a game, *how* can a mediator/adminstrator, who can not enforce behavior, lead the agents to adopt a desired behavior?

The literature about point (1) is almost empty; the celebrated Nash equilibrium can be considered as a partial solution to point (2); however, the guarantees offered by Nash equilibrium are very weak, which leaves that problem without a satisfactory answer. The reason for the lack of work on these subjects is not an accident; the task of coming up with *game-theoretic recommendations* is an extremely challenging one; one may be even tempted to say that addressing these fundamental issues is hopeless. In this abstract we provide some pointers to some of our work in the recent years that provides some progress in this uphill battle. The aim of this short abstract is to report briefly about this progress; the interested reader should consult the corresponding papers in order to obtain further understanding.

2. COMPETITIVE SAFETY ANALYSIS

One of the central challenges of game theory is that of providing a decision maker with an advice about how he should choose his action in a given multi-agent encounter. This challenge, which falls under the so-called prescriptive agenda, has been left without a real answer. For example, the celebrated Nash equilibrium (NE), which is the basis for most game-theoretic analysis, suggests that a multi-agent behavior would be considered "rational" if no decision-maker would prefer to deviate from it, assuming the other decision-makers stick to it. However, while this is a very useful concept from a descriptive point of view, it does not address the question of how *should* a particular agent choose his action in a given game. A NE strategy can only be justified by assuming that the other agents are committed to a specific action profile, which is an unreasonably strong assumption regarding their rationality.

Only very few suggestions have been made in order to address the above challenge. One approach is to suggest to the agent a strategy which will be useful against an opponent taken from a particular class (see e.g. [23]). A related idea is to try and learn the opponent model in a repeated interaction in order to optimize behavior against it [9], and the optimization of behavior against stationary opponents [12]. However, what is common to the above approaches is that there are no guarantees to our agent, unless we severely restrict the opponent he may face. In [34] we have suggested an alternative approach, which is referred to as *competitive safety analysis*, motivated by observation made by Aumann in [6]. This approach deals with guarantees the agent can be provided with, as discussed below.

It is well known that in a purely competitive setting, employing a safety level strategy, one that maximizes the agent's expected utility in the worst case, is the only reasonable mode of behavior. For partially cooperative settings, in [34] we justified the use of safety-level strategy by introducing the notion of *C-competitive safety strategy* – a strategy that *guarantees* a payoff which is not less than a factor C of what is obtained in equilibrium. If there exists a C-competitive strategy for small C, then this strategy is a reasonable suggestion for the decision maker! However, the main challenge is whether for interesting contexts, we do have such competitive safety strategies. Surprisingly, our work has shown the usefulness of this approach in two central settings: congestion settings, and ad auctions.

We show that in an extended decentralized load balancing setting a $9/8$-competitive safety strategy exists. This implies that if an agent has to choose among a fast service provider and a slow service provider, where service is split equally among agents who use a service provider, one can either rely on rationality assumptions leading to Nash equilibrium, or use an algorithm that guarantees $8/9$ of the corresponding value without relying on rationality assumptions. We also discuss extensions of this result to more general settings. In particular, we deal with the cases of arbitrary number of

Cite as: Game-Theoretic Recommendations: Some Progress in an Uphill Battle (Invited Paper), Moshe Tennenholtz, Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008), Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008, Estoril, Portugal, pp. XXX-XXX.

service providers, and arbitrary different speeds of service. We show that a ratio of 4/3 can be obtained when we allow arbitrary speeds for two service providers. We also consider the notion of a k-regular network, where k is the ratio between the average speed of service and the lowest speed of service (by a given set of service providers), and show that a k-competitive safety strategy exists for general k-regular networks.

In [17] we applied competitive safety analysis to the model of ad auctions, which are mechanisms for assigning online advertisement space to agents according to their (proclaimed) utility from using it. The formal model that we use is based on [36] (with minor changes). Needless to say that if useful C-competitive safety strategies exist (i.e. ones that guarantee a relatively small constant factor C) in the ad auction setting, then they may provide useful means for bidders in such auctions.

The basic model of position auctions assumes that bidders' valuations for ad slots are common knowledge. In a more realistic model, each agent knows only his own valuation, while the valuations of all agents are assumed to be selected from some known distribution. We provide an analysis of competitive safety strategies for both the complete and the incomplete information settings. Interestingly, we obtain sharp difference between the usefulness of the approach in the complete and the incomplete information settings. While in the complete information setting, it turns out that no useful competitive safety strategy exists, such strategies exist in the (more realistic) incomplete information setting! Namely, we show that in the complete information setting, assuming N bidders, the value is $O(N)$ with N being almost a tight bound. On the other hand, we show positive results in the incomplete information setting. We consider valuations which are taken from the uniform distribution, and consider two classical types of click-rate functions: the exponential and he linear click-rate functions. We show the existence of an e-competitive strategy for the case of exponential click-rate functions, and a $2 \cdot \frac{N}{N-1}$-competitive safety strategies for linear click-rate functions.

3. ACTION PREDICTION IN ENSEMBLES OF GAMES

While competitive safety analysis is a general approach which perform well in guaranteeing relatively high payoff in some interesting cases, it does not always apply. A general solution in this case may be: try to predict the other agents' behavior and best respond to that. While this may sound naive at first, in a line of research initiated in [1] we show that this approach can indeed lead to highly useful results, by considering *game ensembles*. The idea is that an agent may predict an opponent behavior in a given game in a very effective manner, based on his actions in **other** games, as well other agents' behaviors in other games and the game under consideration. In a sense, what we offer is to try and learn association rules among games, ones that will allow to improve upon prediction of an agent's action in a given game. One of our main contributions is by suggesting this learning approach, and proving its surprising feasibility using real data obtained in experiments involving human subjects. This is done by comparing to existing approaches in the experimental economics and cognitive psychology literature.

The economics literature refers to population learning[33]. In population learning we aim at predicting an agent's action based on statistics on how the population played a game in the past. Population learning is considered a good predictor for an agent's behavior in a game, and therefore will be used as a benchmark.

We chose two leading lines of research to serve as representatives of the modelling approach. The first line of research models agents according to levels of reasoning. The model was proposed by Stahl [31] and Wilson [32]. This approach classifies players into types, based on the number of reasoning steps they appear to do. In a simple variant of this model "level 0" players choose uniformly among the game's strategies, and "level k" players choose the best response to "level $(k-1)$" choices. The roots of the thinking steps approach can be found in Harsanyi's "tracing procedure" [14].

The specific model we compare to, from the above line of research is more sophisticated; it was published in [8] — a Poisson cognitive hierarchy. The cognitive hierarchy theory assumes that players use different numbers of steps of strategic thinking when playing a new one-shot game. To make the theory's predictions precise, players who do k steps of thinking believe others are doing fewer steps, and they guess accurately the relative proportions of less-sophisticated players. This creates a clean model that can be solved inductively: First figure out what 0-step players do, then what 1-step players do (anticipating what 0-step players do), etc. The frequency of k-step thinkers is assumed to be Poisson.

The second central line of research in agent modelling we compare to, characterizes players according to behavioral decision rules [10]. Namely, in this approach Costa Gomes et al. consider two classes of decision rules: strategic and non-strategic, and assign probabilities to each of these decision rules based on empirical data.

The above are strong representatives of the agent modelling approach. In our work we exploit a simple machine learning technique in order to offer an approach for addressing the challenge of action prediction in ensembles of games. In a sense, our work is related to the relatively recent work on case-based decision making[13]. In case-based decision making, the idea of case-based reasoning, which is a classical topic in AI, is exploited to introduce an alternative approach to decision making in strategic contexts. This approach is based on similarity measures between different decision problems. Our work suggests to learn such similarity/association between decision problems, in order to improve upon opponent prediction in games, ultimately improving payoffs. We show that this machine learning approach is highly useful for that context. We have also experimented with other machine learning techniques (such as ID3 and KNN), but they were found to be less efficient for our objectives.

In order to evaluate the different methods meaningful scoring rules should be used as evaluation criteria. Scoring rules provide summary measures for the evaluation of probabilistic forecasts. Two evaluation criteria were taken from Camerer and Costa-Gomes' work - MSE (mean square error) and MLE (maximum likelihood estimator). In addition, we suggest two additional evaluation criteria - absolute prediction and best response. The absolute prediction score is a special case of a zero-one scoring rule. It rewards a probabilistic forecast if the mode of the predictive distribution materializes. In case of multiple modes, the reward is reduced proportionally. This score captures the number of times our prediction was 'right', and is popular among machine learning researchers. The "best response" criterion measures the payoffs one gets by choosing best-response to the corresponding prediction. This is perhaps the criterion which has the highest economic value since it quantifies how much can an agent gain from having the prediction rule at hand.

In our work we gathered data from the founders of the existing modelling techniques, as well as conducted our own experiments with close to 100 human subjects. We show that the newly proposed machine learning technique outperforms the other approaches on all data sets and under all criteria. This gives great hope to the applicability of this technique in equipping an agent with an advice about how to choose its action.

4. MEDIATORS

One of the most basic questions of game theory is: given a game in strategic form, what is the solution of the game? Basically, by a "solution" we mean a *stable* strategy profile which can be proposed to all agents, and no rational agent would want to deviate from it. Many solution concepts for games have been studied, differing mainly by the assumptions that a rational agent would have to make about the rationality of other agents. The best known solution concept for games is the Nash equilibrium. In order to understand some of the issues let us consider the following formalism.

Let N be the set of players in the game, A_i be the set of actions (strategies) available to player i, A the set of action profiles, and let u_i be player i's utility function, a profile of actions $a \in A$ is a *Nash equilibrium* (NE) if

$$\forall i \in N \; a_i \in br_i(a_{-i})$$

Here, $br_i(a_{-i})$ for $i \in N$, $a_{-i} \in A_{-i}$ denotes

$$\arg\max_{a_i \in A_i} \{u_i(a_i, a_{-i})\}$$

(the set of *best responses* of i to a_{-i}).

There are two basic problems with the Nash equilibrium as a solution concept for games:

Problem 1: A NE guarantees absence of profitable deviations to a player only in the case that all the other players play according to the suggested profile; in the case where even one of the other players deviates, we have no such guarantees. So, the assumption that this concept requires about the rationality of other players is: all the other players will stick to their prescribed strategies. But why should a rational player make that assumption?

The following stability concept takes this problem into account: A profile of actions $a \in A$ is an *equilibrium in weakly dominant strategies* if

$$\forall i \in N, b_{-i} \in A_{-i} \quad a_i \in br_i(b_{-i})$$

The above definition strengthens the concept of NE by taking care of the aforementioned problem: no unilateral deviation *can ever be* beneficial, no matter what other players do; in other words, it requires no assumptions on the rationality of other players.

Problem 2: A NE does not take into account joint deviations by coalitions of players. We usually assume that an individual will deviate from a profile if she has an available strategy that strictly increases her income. In some settings it would be natural to assume also that a group of individuals will deviate if they have an available joint strategy that strictly increases the income of each group member. For example, consider the famous Prisoner's Dilemma game:

	C	D
C	4,4	0,6
D	6,0	1,1

The strategy profile (D, D) is a NE and even an equilibrium in weakly dominant strategies; however, it is not stable in the sense that if both players deviate to (C, C), the income of each one of them will increase. The following stability concept by [5] deals with this problem:

A profile of actions $a \in A$ is a *strong equilibrium* (SE) if

$$\forall S \subseteq N \quad a_S \in br_S(a_{-S})$$

Here, the concept of best response strategy is extended to multiple players as follows: for $S \subseteq N$ and $a_{-S} \in A_{-S}$, $br_S(a_{-S})$ denotes the set of best responses of S to a_{-S}: $br_S(a_{-S}) =$
$\{a_S \in A_S | \forall b_S \in A_S \; \exists i \in S \quad u_i(b_S, a_{-S}) \leq u_i(a_S, a_{-S})\}$

A major problem with the above proposed solutions is that they rarely exist. In order to overcome this problem we suggest the study of *mediators*, as tools for leading agents to desired behaviors. A mediator can not design new games, or enforce behaviors by the agents, but he can make reliable offers. In particular, consider the following type of mediator introduced in the study titled k-implementation [20]. The mediator here is reliable party which has only one source of power: his reliability. He can commit to payments to the different agents, when certain observable outcomes will be reached, and the agents can be sure that they will be paid appropriately. However, he can not punish agents or enforce behaviors. As it turns out these mediators can be extremely useful.

Consider the following simple congestion setting.[1] Assume that there are two agents, 1 and 2, that have to select among two service providers (e.g., machines, communication lines, etc.) One of the service providers, f, is a fast one, while the other, s, is a slower one. We capture this by having an agent obtaining a payoff of 6 when he is the only one that uses f, and a payoff of 4 when he is the only one who uses s. If both agents select the same service provider then its speed of operation decreases by a factor of 2, leading to half the payoff. That is, if both agents use f then each one of them obtains a payoff of 3, while if both agents use s then each one of them obtains 2. In a matrix form, this game is described by the following bimatrix:

[1]Congestion in the context of self-motivated parties is a central topic in the recent CS literature [16, 27, 28], as well as in the game theory literature [26, 19]. This example is used for purposes of illustration only; however, the technique used in this example can be extended to arbitrary complex games, as we will later show.

$M =$

	f	s
f	3, 3	6, 4
s	4, 6	2, 2

Assume that our mediator wish to prevent the agents from using the same service provider (leading to low payoffs for both), while relying only on the idea that agents will use dominant strategies. Then it can do as follows: it can promise to pay agent 1 a value of 10 if both agents will use f, and promise to pay agent 2 a value of 10 if both agents will use s. These promises transform M to the following game:

$M' =$

	f	s
f	13, 3	6, 4
s	4, 6	2, 12

Notice that in M', strategy f is dominant for agent 1, and strategy s is dominant for agent 2. As a result the only rational strategy profile is the one in which agent 1 chooses f and agent 2 chooses s. Hence, the mediator implements one of the desired outcomes. Moreover, given that the strategy profile (f, s) is selected the mediator will have to pay nothing. It has just implemented, **in dominant strategies**, a desired behavior (obtained in one of the Nash equilibria) with zero cost, relying only on its creditability, without modifying the rules of interactions or enforcing any type of behavior! In this case we say that the desired behavior has a 0-implementation. More generally, an outcome has a k-implementation if one can make it obtained using dominant strategies with a cost of k.

It can be shown: an outcome is 0-implementable iff it is a Nash equilibrium.

As we can see the idea of k-implementation is extremely useful in handling problem 1, where we wish to implement desired behaviors in dominant strategies.

Another type of mediators has been introduced in order to deal with stability against group deviations [21]. These mediators are *action mediators*; such a mediator is a reliable entity that can interact with the players and perform on their behalf actions in a given game. However, a mediator can not enforce behavior. Indeed, an agent is free to participate in the game without the help of the mediator. This notion is highly natural in a setting in which there exists some form of reliable party or administrator in place. Indeed, many markets employ very powerful forms of mediators, like brokers or routers in communication networks. In order to illustrate the power of a reliable (action) mediator, consider the prisoners dilemma described above. The unique equilibrium is inefficient; indeed, if both agents deviate from defection to cooperation then both of them will improve their payoffs. Formally, mutual defection is not a strong equilibrium. Consider a reliable mediator who offers the agents the following protocol: if **both** agents agree to use the mediator services then he will perform cooperate on behalf of both agents. However, if only one agent agrees to use his services then he will perform defect on behalf of that agent. Notice that when accepting the mediator's offer the agent is committed to actual behavior as determined by the above protocol. However, there is no way to enforce the agents to accept the suggested protocol, and each agent is free to cooperate or defect without using the mediator's services. Hence, the mediator's protocol generates the following mediated game:

	Mediator	Cooperate	Defect
Mediator	4,4	6,0	1,1
Cooperate	0,6	4,4	0,6
Defect	1,1	6,0	1,1

The mediated game has a most desirable property: in this game there is a strong equilibrium; that is, equilibrium which is stable against deviations by coalitions. In this equilibrium both agents will use the mediator services, which will lead them to a payoff of 4 each! We call a strong equilibrium in a mediated game: a *strong mediated equilibrium.*

Given the general concept of a mediator, it can be proved that mediators can indeed significantly increase the set of economic interactions in which desired outcomes, which are stable against deviations by coalitions, can be obtained. For example, every *balanced* symmetric game possesses a strong mediated equilibrium, which also leads to optimal surplus, where a game in strategic form is called balanced if its associated core (see [37]) is non-empty. Another positive result is with regard to deviations of coalitions of size at most k: any symmetric game with n agents, if $k!$ divides n then there exists a k-strong mediated equilibrium, leading to optimal surplus.[2]. However, if $k!$ does not divide n, then it is shown

[2]As an anecdote, the Parlament in Israel contains $120 = 5!$ members. Hence, every anonymous game played by this Parlament possesses an optimal surplus symmetric 5-strong mediated equilibrium. While no Parlament member is able to give the right of voting to a mediator, this right of voting

that the game may or may not possess a k-strong mediated equilibrium.

At this point one may see that k-implementation and action mediators have been really useful in overcoming central problems, and implementing desired behaviors as dominant strategies and strong equilibrium, respectively. However, one may wish to go even further and address simultaneously both issues. This brings us to the stability introduced in [29]: a profile of actions $a \in A$ is an *equilibrium in group (weakly) dominant strategies* (GDS) if

$$\forall S \subseteq N, b_{-S} \in A_{-S} \quad a_S \in br_S(b_{-S})$$

Existence of a GDS implies, for each player, that no matter what the other players choose, and no matter with whom can she unite in making her decision, they will not find a joint strategy that will be better to all of them than the proposed one. And thus, if a GDS exists in a given game, we can safely declare it to be the solution of the game. However, a GDS does not exist in any game that has ever been a subject of interest. This is not surprising, since the concept is so strong that its mere existence renders any game not interesting. For this reason, the concept was never a subject of exploration in complete information games. In incomplete information games the concept is known under the name of *group strategy proofness* and is widely studied, because in some cases such solutions can be indeed implemented by mechanism design. However, the whole approach of mechanism design is not applicable to complete information games – although we would indeed want to assume the existence of an interested party, we don't want to give it the power to design the game. Recent work [29] showed that by using mediators GDS can be implemented in a very general and natural class of games. The mediators combine the capabilities of k-implementation and an extended version of the action mediators (termed routing mediators [30]).

5. TRUST-BASED RECOMMENDATION SYSTEMS

The previous sections offered surprisingly powerful solutions to the issues raised in the introduction. However, game-theoretic recommendations are not restricted only to non-cooperative games; in particular, they are desired also in settings derived from the fundamental setting of social choice, and its adaptation to handling ranking/trust/reputation in multi-agent systems.

In the classical theory of social choice, a theory developed by game-theorists and theoretical economists, we consider a set of agents (voters) and a set of alternatives. Each agent ranks the alternatives, and the major aim is to find a good way to aggregate the individual preferences into a social preference. The major tool offered in this theory is the axiomatic approach: study properties (termed axioms) that characterize particular aggregation rules, and analyze whether particular desired properties can be simultaneously satisfied. In a ranking system [2] the set of voters and the set of alternatives coincide, e.g. they are both the pages in the web; in this case the links among pages are interpreted as votes: pages that page p links to are preferable by page p to pages it does not link to; the problem of preference aggregation becomes the problem of page ranking. Trust systems are personalized ranking systems [3] where the ranking is done for (and from the perspective of) each individual agent. Here the idea is to see how to rank agents from the perspective of a particular agent/user, based on the trust network generated by the votes. In a trust-based recommendation system the agents also express opinions about external topics, and a user who has not expressed an opinion should be recommended one based on the opinions of others and the trust network [4]. Hence, we get a sequence of very interesting settings, extending upon classical social choice, where the axiomatic approach can be used.

---

may be replaced in real life by a commitment to follow the mediator's algorithm

On the practical side, ranking, reputation, recommendation, and trust systems have become essential ingredients of web-based multi-agent systems (e.g. [15, 25, 7, 35, 11]). These systems aggregate agents' reviews of products and services, and of each other, into valuable information. Notable commercial examples include Amazon and E-Bay's recommendation and reputation systems (e.g. [24]), Google's page ranking system [22], and the Epinions web of trust/reputation system (e.g. [18]). Our work shows that an extremely powerful way for the study and design of such systems is the axiomatic approach, extending upon the classical theory of social choice [35, 2, 3]. Below we discuss some of the details of that work.

Consider a setting where there is a single item of interest (e.g., a product, service, or political candidate). A subset of the agents have prior opinions about this item. Any of the remaining agents might desire to estimate whether or not they would like the item, based on the opinions of others. In the off-line world, a person might first consult her friends for their recommendations. In turn, the friends, if they do not have opinions of their own, may consult their friends, and so on. Based on the cumulative feedback the initial consulter receives, she might form her own subjective opinion. An automated trust-based recommendation system aims to provide a similar process to produce high-quality *personalized* recommendations for agents.

In [4] we model this setting as an annotated directed graph in which some of the nodes are labelled by votes of $+$ and $-$. Here a node represents an agent, and an edge directed from a to b represents the fact that agent a trusts agent b. A subset of the nodes are labelled by $+$ or $-$ votes, indicating that these nodes have already formed opinions about the item under question. Based on this input, a recommendation system must output a recommendation for each unlabelled node. We call such an abstraction a *voting network* because it models a variety of two-candidate voting systems, where the candidates are $+$ and $-$. For an example, consider a directed star graph where a single root node points to n agents with labels, which models a committee making a recommendation to the root node. In that setting, majority and consensus are two common voting rules. For another example, the U.S. presidential voting system can be modelled as a more complicated digraph, where the root points to nodes representing the members of the electoral college, and the electoral college nodes point to nodes representing the voters in the state or congressional district that they represent.

A multitude of recommendation systems have been proposed and implemented, and many fit into the network-based framework described above. This raises the question of how to determine the relative merits of alternative approaches to providing trust-based recommendations. The task of comparing recommendation systems is complicated by the difficulty of producing an objective measure of recommendation quality.

In our work we import the axiomatic approach from the theory of social choice, and use it to compare and contrast recommendation systems. We consider two styles of axiomatic analysis: the descriptive approach, where we aim to find a set of axioms that characterizes a given system; and the normative approach, where we first come up with a set of axioms, and then determine which recommendation systems will satisfy all or some of the axioms in the set.

We begin with an impossibility theorem: for a small, natural set of axioms, there is no recommendation system simultaneously consistent with all axioms in the set. However, for any proper subset of the axioms there exists a recommendation system that satisfies all axioms in the subset. We consider two ways past this negative result, both by replacing the *transitivity* axiom (defined in our work). We prove that there are recommendation systems consistent with both new sets of axioms. We also show that when one of these new sets is augmented with an additional axiom, the resulting set of axioms is uniquely satisfied by a recommendation system based on random walks.

We also consider the descriptive approach, in which we characterize existing (acyclic) systems, like simple committees and the U.S. presidential elections, by a simple majority axiom. We generalize this to an axiom that leads to a unique "minimum cut" system on general undirected (possibly cyclic) graphs.

We define a notion of incentive compatibility for recommendation systems. This is important when designing systems for deployment in monetized settings, because, as experience has shown, self-interested agents will not respect the rules of the system when there is money to be made by doing otherwise. We find that all of the recommendation systems for which we provide a characterizing set of axioms turn out to be incentive compatible, including the random walk system, majority of majorities system, and minimum cut system. In contrast, the personalized PageRank system and various other natural systems are not incentive compatible.

6. CONCLUSIONS

Our work considers two highly challenging problems in the foundations of game theory and its application to multi-agent systems. Namely, we consider the question of how should an agent choose its action in a given game, and the task of leading agents to adopt desired behaviors in a given game. In the recent years we provided some useful attacks on these fundamental problems. Our studies of competitive safety analysis and the study of learning in ensembles of games, provide surprisingly useful tools in addressing the first challenge. Our theory of mediators provides powerful tools in addressing the second challenge. These approaches refer to non-cooperative games; in the context of social choice, we provide a complementary work by proposing the axiomatic approach to ranking/reputation/trust/recommendation systems; in particular, our work on trust-based recommendation systems introduce several basic results in the characterization of desired recommendation techniques in that context.

7. REFERENCES

[1] A. Altman, A. Boden-Bercovici, and M. Tennenholtz. Learning in one-shot strategic form games. In *Proceedings of ECML-06*, 2006.

[2] A. Altman and M. Tennenholtz. On the axiomatic foundations of ranking systems. In *Proc. 19th International Joint Conference on Artificial Intelligence*, pages 917–922, 2005.

[3] A. Altman and M. Tennenholtz. An axiomatic approach to personalized ranking systems. In *Proc. 20th International Joint Conference on Artificial Intelligence*, 2006.

[4] R. Andersen, C. Borgs, J. Chayes, U. Feige, A. Flaxman, A. Kalai, V. Mirrokni, and M. Tennenholtz. Trust-Based Recommendation Systems: an Axiomatic Approach. In *Proceedings of WWW-08*, 2008.

[5] R. Aumann. Acceptable points in general cooperative n-person games. In A. Tucker and R. Luce, editors, *Contribution to the Theory of Games, Vol. IV, Annals of Mathematics Studies, 40*, pages 287–324. 1959.

[6] R. Aumann. On the non-transferable utility value: A comment on the Roth-Shaper examples. *Econometrica*, 53(3):667–677, 1985.

[7] Y. Bakos and C. N. Dellarocas. Cooperation without enforcement? a comparative analysis of litigation and online reputation as quality assurance mechanisms. MIT Sloan School of Management Working Paper No. 4295-03, 2003.

[8] C. Camerer, T. Ho, and J. Chong. A cognitive hierarchy model of games. *The Quarterly Journal of Economics*, 119(3):861–898, 2004.

[9] D. Carmel and S. Markovitch. Exploration strategies for model-based learning in multiagent systems. *Autonomous Agents and Multi-agent Systems*, 2(2):141–172, 1999.

[10] M. Costa-Gomes, V. P. Crawford, and B. Broseta. Cognition and behavior in normal-form games: An experimental study. *Econometrica*, 69(5):1193–1235, 2001.

[11] R. Dash, S. Ramchurn, and N. Jennings. Trust-based mechanism design. In *Proceedings of the Third International Joint Conference on Autonomous Agents and MultiAgent Systems*, pages 748–755, 2004.

[12] Y. Freund and R. E. Schapire. Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29:79–103, 1999.

[13] I. Gilboa and D. Schmeidler. Case-based decision theory. *Quarterly Journal of Economics*, 110:605–639, 1995.

[14] R. S. John C. Harsanyi. A general theory of equilibrium selection in games. *Journal of Economic Literature*, 27:1171–1173, 1989.

[15] J. M. Kleinberg. Authoritative sources in a hyperlinked environment. *Journal of the ACM (JACM)*, 46(5):604–632, 1999.

[16] E. Koutsoupias and C. Papadimitriou. Worst-Case Equilibria. In *STACS*, 1999.

[17] D. Kuminov, , and M. Tennenholtz. Competitive Safety Analysis in Position Auctions. In *Proceedings of WINE-07*, 2007.

[18] P. Massa and P. Avesani. Controversial users demand local trust metrics: An experimental study on epinions.com community. In *Proc. of AAAI-05*, pages 121–126, 2005.

[19] D. Monderer and L. Shapley. Potential games. *Games and Economic Behavior*, 14:124–143, 1996.

[20] D. Monderer and M. Tennenholtz. K-Implementation. *Journal of Artificial Intelligence Research*, 21:37–62, 2004.

[21] D. Monderer and M. Tennenholtz. Strong Mediated Equilibrium. In *Proceedings of AAAI-06*, 2006.

[22] L. Page, S. Brin, R. Motwani, and T. Winograd. The

pagerank citation ranking: Bringing order to the web. Technical Report, Stanford University, 1998.

[23] R. Powers and Y. Shoham. New criteria and a new algorithm for learning in multi-agent systems. Proceedings of NIPS 2004, 2004.

[24] P. Resnick and R. Zeckhauser. Trust among strangers in internet transactions: Empirical analysis of ebay's reputation system. Working Paper for the NBER workshop on empirical studies of electronic commerce, 2001.

[25] P. Resnick, R. Zeckhauser, R. Friedman, and E. Kuwabara. Reputation systems. *Communications of the ACM*, 43(12):45–48, 2000.

[26] R. Rosenthal. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.

[27] T. Roughgarden. The price of anarchy is independent of the network topology. In *Proceedings of the 34th Annual ACM Symposium on the Theory of Computing*, pages 428–437, 2001.

[28] T. Roughgarden and E. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, April 2002.

[29] O. Rozenfeld and M. Tennenholtz. Group Dominant Strategies. In *Proceedings of WINE-07*, 2007.

[30] O. Rozenfeld and M. Tennenholtz. Routing Mediators. In *Proc. 20th International Joint Conference on Artificial Intelligence*, 2007.

[31] D. Stahl. Evolution of smart players. *Games and Economic Behavior*, 5(4):604–617, October 1993.

[32] D. I. Stahl and P. W. Wilson. Experimental evidence on players' models of other players. *Journal of Economic Behavior & Organization*, 25(3):309–327, December 1994.

[33] D. O. Stahl. Population rule learning in symmetric normal-form games: theory and evidence. *Journal of Economic Behavior and Organization*, 1304:1–14, 2001.

[34] M. Tennenholtz. Competive Safety Analyis: robust decision-making in multi-agent systems. *Journal of Artificial Intelligence Research*, 17:363–378, 2002.

[35] M. Tennenholtz. Reputation systems: An axiomatic approach. Proceedings of UAI-04, 2004.

[36] H. R. Varian. Position auctions. Technical Report, UC Berkeley. Available at http://www.sims.berkeley.edu/∼hal/Papers/2006/, 2006.

[37] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, 1944.

Moshe Tennenholtz is a professor with the faculty of Industrial Engineering and Management at the Technion--Israel Institute if technology, where he holds the Sondheimer Technion Academic Chair. During 1999-2002 he has been a visiting professor at Stanford CS department, where he has also been a research associate during the years 1991-1993. Moshe received his B.Sc. in Mathematics from Tel-Aviv University (1986), and his M.Sc. and Ph.D. (1987, 1991) from the Department of Applied Mathematics and Computer Science in the Weizmann Institute. Moshe served as the editor-in-chief of the Journal of Artificial Intelligence Research [JAIR]; he is an associate editor of Games and Economic Behavior, and of the international journal of autonomous agents and multi-agent systems, and an editorial board member of the Journal of Machine Learning Research [JMLR]. In joint work with colleagues and students he introduced several contributions to the interplay between computer science and game theory, such as the study of artificial social systems, co-learning, non-cooperative computing, distributed games, the axiomatic approach to qualitative decision making, the axiomatic approach to ranking, reputation, and trust systems, competitive safety analysis, program equilibrium, mediated equilibrium, and learning equilibrium.