

Competitive Safety Analysis

Moshe Tennenholtz

Computer Science Department
Stanford University
Stanford, CA 94305 *

Abstract

Much work in AI deals with the selection of proper actions in a given (known or unknown) environment. However, the way to select a proper action when facing other agents is quite unclear. Most work in AI adopts classical game-theoretic equilibrium analysis to predict agent behavior in such settings. Needless to say, this approach does not provide us with any guarantee for the agent. In this paper we introduce competitive safety analysis. This approach bridges the gap between the desired normative AI approach, where a strategy should be selected in order to guarantee a desired payoff, and equilibrium analysis. We show that a safety level strategy is able to guarantee the value obtained in a Nash equilibrium, in several classical computer science settings. Then, we discuss the concept of competitive safety strategies, and illustrate its use in a decentralized load balancing setting, typical to network problems. In particular, we show that when we have many agents, it is possible to guarantee an expected payoff which is a factor of $8/9$ of the payoff obtained in a Nash equilibrium. Finally, we discuss the extension of the above concepts to Bayesian games, and illustrate their use in a basic auctions setup.

Introduction

Deriving solution concepts for multi-agent encounters is a major challenge for researchers in various disciplines. The most famous and popular solution concept in the economics literature is the Nash equilibrium. Although Nash equilibrium and its extensions and modifications are powerful descriptive tools, and have been widely used in the AI literature (see e.g. (Rosenschein & Zlotkin 1994; Kraus 1997; Sandholm & Lesser 1995)), their appeal from a normative AI perspective is somewhat less satisfactory.¹ We wish

*Permanent address: Faculty of Industrial Engineering and Management, Technion, Haifa 32000, Israel
Copyright © 2002, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

¹If we restrict ourselves to cases where there exists an equilibrium in dominant strategies (as is done in some of the CS literature; see e.g. (Nisan & Ronen 1999)) then the corresponding equilibrium is appealing from a normative perspective. However, such cases rarely exist.

to equip an agent with an action that guarantees some desired outcome, or expected utility, without relying on other agents' rationality.

This paper shows that, surprisingly, the desire for obtaining a guaranteed expected payoff, where this payoff is of the order of the value obtained in a Nash equilibrium, is achievable in various classical computer science settings. Our results are inspired by several interesting examples for counter-intuitive behaviors obtained by following Nash equilibria and other solution concepts (Roth 1980; Aumann 1985). One of the most interesting and challenging examples has been introduced by Aumann (Aumann 1985). Aumann presented a 2-person 2-choice (2×2) game g , where the safety-level (probabilistic maximin) strategy of the game is not a Nash equilibrium of it, but it does yield the expected payoff of a Nash equilibrium of g . This observation may have significant positive ramifications from an agent's design perspective. If a safety-level strategy of an agent guarantees an expected payoff that equals its expected payoff in a Nash equilibrium, then it can serve as a desirable robust protocol for the agent!

Given the above, we are interested in whether an optimal safety level strategy leads to an expected payoff similar to the one obtained in a Nash equilibrium of simple games that represent basic variants of classical computer science problems. As we show, this is indeed the case for 2×2 games capturing simple variants of the classical load balancing and leader election problems.

A more general question refers to more general 2×2 games. We show that if the safety-level strategy is a (strictly) mixed one, then its expected payoff is identical to the expected payoff obtained in a Nash equilibrium in any generic non-reducible 2×2 game. We also show that this is no longer necessarily the case if we have a pure safety-level strategy. In addition, we consider general 2-person set-theoretic games (which naturally extend 2×2 leader election games) and show that if a set-theoretic game g possesses a strictly mixed strategy equilibrium then the safety level value for a player in that game equals the expected payoff it obtains in that equilibrium.

Following this, we define the concept of C -competitive safety strategies. Roughly speaking, a strategy will be called

a C -competitive safety strategy, if it guarantees an expected payoff that is $\frac{1}{C}$ of the expected payoff obtained in a Nash equilibrium. We show that in an extended decentralized load balancing setting a $9/8$ -competitive strategy exists, when the number of players is large. We also discuss extensions of this result to more general settings. Then, we discuss C -competitive strategies in the context of Bayesian games. In particular we show the existence of an e -competitive safety strategy for a classical first-price auctions setup.

Previous work has been concerned with comparing the payoffs that can be obtained by an optimal centralized (and Pareto-efficient) controller to the expected payoffs obtained in the Nash-equilibria of the corresponding game (Koutsoupias & Papadimitriou 1999). That work is in the spirit of competitive analysis, a central topic in theoretical computer science (Borodin & El-Yaniv 1998). Our work can be considered as suggesting a complementary approach, comparing the safety-level value to the agent's expected payoff in a Nash equilibrium. Needless to say that in computational settings, where failures are possible, and rationality assumptions about participants' behavior should be minimized, a safety-level strategy has a special appeal, especially when it yields a value that is close to the expected payoff obtained in a Nash equilibrium.

Basic definitions and notations

A *game* is a tuple $G = \langle N, \{S_i\}_{i=1}^n, \{U_i\}_{i=1}^n \rangle$, where N is a set of n players, S_i is a finite set of pure strategies available to player i , and $U_i : \prod_{i=1}^n S_i \rightarrow \mathbb{R}$ is the payoff function of player i .

Given S_i , we denote the set of probability distributions over the elements of S_i by $\Delta(S_i)$. An element $t \in \Delta(S_i)$ is called a *mixed strategy* of player i . It is called a *pure strategy* if it assigns probability 1 to an element of S_i , and it is called a *strictly mixed strategy* if it assigns a positive probability to each element in S_i . A tuple $t = (t_1, \dots, t_n) \in \prod_{i=1}^n \Delta(S_i)$ is called a *strategy profile*. We denote by $U_i(t)$ the expected payoff of player i given the strategy profile t .

A strategy profile $t = (t_1, \dots, t_n)$ is a *Nash equilibrium* if $\forall i \in N, U_i(t) \geq U_i(t_1, \dots, t_{i-1}, t'_i, t_{i+1}, \dots, t_n)$ for every $t'_i \in S_i$. The Nash equilibrium $t = (t_1, \dots, t_n)$ is called a *pure strategy Nash equilibrium* if t_i is a pure strategy for every $i \in N$. The Nash equilibrium $t = (t_1, \dots, t_n)$ is called a *strictly mixed strategy Nash equilibrium* if for every $i \in N$ we have that t_i is a strictly mixed strategy.

Given a game g and a mixed strategy of player i , $t \in \Delta(S_i)$, the safety level value obtained by i when choosing t in the game g , denoted by $val(t, i, g)$, is the minimal expected payoff that player i may obtain when employing t against arbitrary strategy profiles of the other players. A strategy t' of player i for which $val(\cdot, i, g)$ is maximal is called a *safety-level strategy* (or a probabilistic maximin strategy) of player i .

A strategy $e \in S_i$ *dominates* a strategy $f \in S_i$ if for every $(s_1, s_2, \dots, s_{i-1}, s_{i+1}, \dots, s_n) \in \prod_{j \neq i} \Delta(S_j)$ we have

$U_i(s_1, \dots, s_{j-1}, e, s_{j+1}, \dots, s_n) \geq U_i(s_1, \dots, s_{j-1}, f, s_{j+1}, \dots, s_n)$, with a strict inequality for at least one such tuple.

A game is called *non-reducible* if there do not exist $e, f \in S_i$, for some $i \in N$, such that e dominates f . A game is called *generic* if for every $i \in N$, and pair of strategies $e, f \in \prod_{j=1}^n S_j$, we have that $U_i(e) = U_i(f)$ only if player i 's strategies in e and f coincide.

A game is called a 2×2 game if $n = 2$ and $|S_1| = |S_2| = 2$.

Decentralized load balancing

In this section we consider decentralized load balancing, where two rational players need to submit messages in a simple communication network: a network of two parallel communication lines e_1, e_2 connecting nodes s and t . Each player has a message that he needs to deliver from s to t , and he needs to decide on the route to be taken. The communication line e_1 is a faster one, and therefore the value of transmitting a single message along e_1 is $X > 0$ while the value of transmitting a single message along e_2 is αX for some $0.5 < \alpha < 1$. Each player needs to decide on the communication line to be used for sending its message from s to t . If both players choose the same communication line then the value for each one of them drops in a factor of two (a player will obtain $\frac{X}{2}$ if both players choose e_1 , and a player will obtain $\frac{\alpha X}{2}$ if both players choose e_2). In a matrix form, this game can be presented as follows:

$$M = \begin{pmatrix} X/2, X/2 & X, \alpha X \\ \alpha X, X & \alpha X/2, \alpha X/2 \end{pmatrix}$$

Proposition 1 *The optimal safety-level value for a player in the decentralized load balancing game equals its expected payoff in the strictly mixed strategy equilibrium of that game.*

The proof of the above proposition appears in the full paper. The above proposition shows that an agent can *guarantee* itself an expected payoff that equals its payoff in a Nash equilibrium of the decentralized load balancing game. This is obtained using a strategy that differs from the agent's strategies in the Nash equilibria of that game (which do not provide that guarantee).

Leader election: decentralized voting

In a leader election setting, the players vote about the identity of the player who will take the lead on a particular task. A failure to obtain agreement about the leader is a bad output, and can be modelled as leading to a 0 payoff. Assume that the players' strategies are either "vote for 1" or "vote for 2", denoted by a_1, a_2 respectively, then $U_i(a_j, a_k) > 0$, where $i, j, k \in \{1, 2\}$, and $j \neq k$. Notice that this setting captures various forms of leader election, e.g. when a player

prefers to be selected, when it prefers the other player to be selected, etc. In a matrix form, this game can be presented as follows (where $a, b, c, d > 0$):

$$M = \begin{pmatrix} a, b & 0, 0 \\ 0, 0 & c, d \end{pmatrix}$$

Proposition 2 *The optimal safety-level value for a player in the leader election game equals its expected payoff in the strictly mixed strategy equilibrium of that game.*

The proof of the above proposition appears in the full paper. The above proposition shows that a agent can *guarantee* itself an expected payoff that equals its payoff in a Nash equilibrium of the leader election game. As in the decentralized load balancing game, this is obtained using a strategy that differs from the agent's strategies in the Nash equilibria of that game (which do not provide that guarantee).

Safety level in general 2×2 games

The results presented in the previous sections refer to 2-person 2-choice variants of central problems occurring in computational contexts. However, it is of interest to see whether these can be extended to other forms of 2×2 games. It is easy to observe that both the load balancing and the leader election settings can be represented as non-reducible generic 2×2 games. The same is true with regard to the game presented by Aumann:

$$M = \begin{pmatrix} 2, 6 & 4, 2 \\ 6, 0 & 0, 4 \end{pmatrix}$$

We can now show:

Theorem 1 *Let G be a 2×2 non-reducible generic game. Assume that the optimal safety level value of a player is obtained by a strictly mixed strategy, then this value coincides with the expected payoff of that player in a Nash equilibrium of G .*

Sketch of proof: Denote the strategies available to the players by a_1, a_2 . Use the following notation: $a = U_1(a_1, a_1), b = U_1(a_1, a_2), c = U_1(a_2, a_1), d = U_1(a_2, a_2), e = U_2(a_1, a_1), f = U_2(a_1, a_2), g = U_2(a_2, a_1), h = U_2(a_2, a_2)$

If a strictly mixed strategy Nash equilibrium exists then it should satisfy that:

$$qa + (1 - q)b = qc + (1 - q)d$$

and

$$pe + (1 - p)g = pf + (1 - p)h$$

where p and q are the probabilities for choosing a_1 by players 1 and 2, respectively.

We get that we should have $qa + b - qb = qc + d - qd$, which implies that $q(a - b - c + d) = d - b$. Similarly, we get that we should have $pe + g - pg = pf + h - ph$, which implies that $p(e - g - f + h) = h - g$.

Hence, in a strictly mixed strategy Nash equilibrium we should have:

$$q = \frac{d - b}{a - b - c + d}$$

and

$$p = \frac{h - g}{e - g - f + h}$$

Notice that since the game is generic then $d \neq b$. If $d > b$ then if q is not in between 0 and 1 then $c > a$ which will contradict non-reducibility. If $d < b$ then in if q is not in between 0 and 1 then $a > c$, which also contradicts non-reducibility. Similarly, since the game is generic then $h \neq g$. If $h > g$ then if p is not in between 0 and 1 then $f > e$ which will contradict non-reducibility. If $h < g$ then in if p is not in between 0 and 1 then $e < f$, which also contradicts non-reducibility.

Given the above we get that p and q define a strictly mixed strategy equilibrium of G .

Consider now the safety level strategy of player 1. If player 1 chooses a_1 with probability p' then it satisfies that:

$$p'a + (1 - p')c = p'b + (1 - p')d$$

This implies that we need to have $p'a + c - p'c = p'b + d - p'd$, which implies $p'(a - c - b + d) = d - c$. Hence, we have

$$p' = \frac{d - c}{a - c - b + d}$$

and

$$1 - p' = \frac{a - b}{a - c - b + d}$$

Compute now the expected payoff for player 1 in the strictly mixed Nash equilibrium, given that $1 - q = \frac{a - c}{a - b - c + d}$, we have that:

$$qa + (1 - q)b = \frac{(d - b)a + (a - c)b}{a - b - c + d} = \frac{da - cb}{a - b - c + d}$$

The expected payoff of the safety level strategy for player 1 will be:

$$p'a + (1 - p')c = \frac{(d - c)a + (a - b)c}{a - b - c + d} = \frac{da - cb}{a - b - c + d}$$

Hence, we get that the expected payoffs of the Nash equilibrium and the safety level strategies for player 1 coincide. The computation for player 2 is similar. ■

The case of pure safety-level strategies

The reader may wonder whether the previous result can be also proved for the case where there are no restrictions on the structure of the safety-level strategy of the game g . As we now show, there exists a generic non-reducible 2×2 game g , where the optimal safety level strategy for a player is pure, and the expected payoff for that player is lower than the expected payoff for that player in all Nash equilibria of g .

Consider a game g , where $U_1(1, 1) = 100, U_1(1, 2) = 40, U_1(2, 1) = 60, U_1(2, 2) = 50$, and $U_2(1, 1) = 100, U_2(1, 2) = 210, U_2(2, 1) = 200, U_2(2, 2) = 90$. In a matrix form this game looks as follows:

$$M = \begin{pmatrix} 100, 100 & 40, 210 \\ 60, 200 & 50, 90 \end{pmatrix}$$

It is easy to check that g is generic and non-reducible. The game has no pure Nash equilibria. In a strictly mixed strategy equilibrium the probability q of choosing a_1 by player 2 should satisfy $100q + 40(1 - q) = 60q + 50(1 - q)$, i.e. that $60q + 40 = 10q + 50$, $q = 0.2$. In that equilibrium the probability that player 1 will choose a_1 is $p = 0.5$, and the expected payoff of player 1 is $100q + 40(1 - q) = 52$. The safety-level strategy for player 1 is to perform a_2 , guaranteeing a payoff of 50, given that (a_2, a_2) is a saddle point in a zero-sum game where the payoffs of player 2 are taken to be the complement to 0 of player 1's original payoffs. Hence, the value of the safety level strategy for player 1 is $50 < 52$.

■

Beyond 2×2 games

The leader election game is an instance of a more general set of games: *set-theoretic games*. In a set theoretic game the sets of strategies available to the players are identical, and the payoff of each player is uniquely determined by the *set* of strategies selected by each player. For example, in a 2-person set-theoretic game we will have that $U_1(s, t) = U_1(t, s), U_2(s, t) = U_2(t, s)$ for every $s, t \in S_1 = S_2$. Notice that set-theoretic games are very typical to voting contexts.

We can prove the following:

Proposition 3 *Given a 2-person set theoretic game g with a strictly mixed strategy Nash equilibrium, then the value of an optimal safety level strategy of a player equals its expected payoff in that equilibrium.*

The proof of the above proposition appears in the full paper. Notice that the proposition considers games with a strictly mixed strategy Nash equilibrium. The proposition does not hold without this restriction.

Competitive safety strategies

Let S be a set of strategies. Consider a sequence of games $(g_1, g_2, \dots, g_j, \dots)$ where i is a player at each of them, its set of strategies at each of these games is S , and there are j players, in addition to i , in g_j . As an example, consider a sequence of decentralized load balancing settings. The $(n-1)$ -th game in this extended load-balancing setting will consist of n players, one of them is i . The players submit their messages along e_1 and e_2 . The payoff for player i when participating in an n -person decentralized load balancing game is $\frac{X}{k}$ (resp. $\frac{\alpha X}{k}$) if he has chosen e_1 (resp. e_2) and additional $k-1$ participants have chosen that communication line.

A mixed strategy $t \in \Delta(S)$ will be called a *C-competitive safety strategy* if there exists some constant $C > 0$, such that

$$\lim_{j \rightarrow \infty} \frac{\text{nash}(i, g_j)}{\text{val}(t, i, g_j)} \leq C$$

where $\text{nash}(i, g_j)$ is the lowest expected payoff player i might obtain in some equilibrium of g_j , and $\text{val}(t, i, g_j)$ is the expected payoff guaranteed for i by choosing t in the game g_j .

The extended decentralized load balancing setting is a typical and basic network problem. If C is small, a *C-competitive safety strategy* for that context will provide a useful protocol of behavior.

We can show:

Theorem 2 *There exists a $9/8$ -competitive safety strategy for the extended decentralized load-balancing setting.*

Sketch of proof: Consider the following strategy profile for the players in an n -person decentralized load balancing game: players $\{1, 2, \dots, \lceil \frac{1}{1+\alpha} n \rceil\}$ will choose e_1 , and the rest will choose e_2 . W.l.o.g we assume that $i = 1$ is the player for which we will make the computation of expected payoffs. It is easy to verify that the above strategy profile is an equilibrium of the game, with an expected payoff for player i that is bounded above by

$$\frac{X(1 + \alpha)}{n} \quad (*)$$

Consider now the following strategy t for player i : select e_1 with probability $\frac{\alpha}{1+\alpha}$ and select e_2 with probability $\frac{1}{1+\alpha}$. It is easy to see that t (if adopted by all participants) is not a Nash equilibrium. However, we will show that it is a competitive safety strategy for small $C > 0$. Consider an arbitrary number of participants n , where $\beta(n-1)$ of the other (i.e. excluding player i) $n-1$ participants use e_2 while the rest use e_1 , for some arbitrary $0 \leq \beta \leq 1$. The expected payoff obtained using t will be:

$$\frac{1}{1 + \alpha \beta(n-1) + 1} + \frac{\alpha}{1 + \alpha} \frac{X}{(1 - \beta)(n-1) + 1}$$

This value is greater or equal to:

$$\frac{1}{1+\alpha} \frac{\alpha X}{\beta n + 1} + \frac{\alpha}{1+\alpha} \frac{X}{(1-\beta)n + 1}$$

The above equals

$$\frac{X\alpha}{1+\alpha} \left[\frac{1}{\beta n + 1} + \frac{1}{(1-\beta)n + 1} \right]$$

Simplifying the above we get:

$$\frac{X\alpha}{1+\alpha} \frac{n+2}{(1+\beta n)(n-\beta n+1)} = (**)$$

Dividing (*) by (**) we get that the ratio is:

$$\frac{(1+\alpha)^2}{\alpha} \frac{(\beta - \beta^2)n^2 + n + 1}{n(n+2)}$$

When n approaches infinity the above ratio approaches

$$\frac{(1+\alpha)^2}{\alpha} (\beta - \beta^2)$$

Given that $0.5 \leq \alpha < 1$ and $0 \leq \beta \leq 1$ we get that the above ratio is bounded by $9/8$ as desired. ■

Extensions: Arbitrary speeds and m links

In this section we generalize the result obtained in the context of decentralized load balancing to the case where we have m parallel communication lines leading from source to target. The value obtained by the agent (w.l.o.g. agent 1) when submitting its message along line i , where n_i agents have decided to submit their messages through that line is given by $\frac{X \cdot \alpha_i}{n_i}$, where $1 = \alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_m > 0$.

Our extension will enable to handle the general binary case where $0 < \alpha < 1$, as well as to discuss cases where a safety level strategy can be very effective in the general m -lines situation.

Using the ideas developed for the case $m = 2$, we can now show:

Theorem 3 *There exists a $\frac{\sum_{i=1}^m \alpha_i \sum_{j=1}^m \prod_{j \neq i} \alpha_j}{m^2 \prod_{i=1}^m \alpha_j}$ -competitive safety strategy for the extended decentralized load-balancing setting, when we allow m (rather than only 2) parallel communication lines, and arbitrary α_i 's.*

The proof of Theorem 3, as well as of the following corollaries appear in the full paper.

In the general binary case, where $\alpha_1 = 1$, and $\alpha_2 = \alpha$, where $0 < \alpha \leq 1$, the above implies the existence of an

$$\frac{(1+\alpha)^2}{4\alpha}$$

competitive strategy.

Corollary 1 *Given an extended load balancing setting, where $m = 2$, with arbitrary speeds of the communication lines ($0 < \alpha \leq 1$), there exists a $\frac{4}{3}$ -competitive strategy.*

Consider now the general m -links (i.e. m parallel communication lines) case. The average network quality (or speed), Q , can be defined as $\frac{\sum_{i=1}^m \alpha_i}{m}$. A network will be called k -regular if $\frac{Q}{\alpha_m} \leq k$. Many networks are k -regular for small k . For example, if $\alpha_m \geq 0.5$ as before, then the network is 2-regular regardless of the number of edges.

Corollary 2 *Given a k -regular network, there exists a k -competitive safety strategy for the extended decentralized load-balancing setting, when we allow m (rather than only 2) parallel edges.*

Together, Theorem 3 and corollaries 1 and 2 extend the results on decentralized load balancing to the general case of m parallel communication lines.

Competitive safety analysis in Bayesian games

The results presented in the previous sections refer to games with complete information. Similar ideas however can be applied to games with incomplete information. We now show the use of competitive safety strategies in games with incomplete information. We have chosen to consider a very basic mechanism, the first-price auction. The selection of first-price auction is not an accident. Auctions are fundamental to the theory of economic mechanism design, and among the auctions that do not possess a dominant strategy, assuming the independent private value model, first-price auctions are probably the most common ones.

We consider a setting where a good g is put for sale, and there are n potential buyers. Each such buyer has a valuation (i.e. maximal willingness to pay) for g that is drawn from a uniform distribution on the interval of real numbers $[0, 1]$. The valuations are assumed to be independent from one another. In a first price auction, each potential buyer is asked to submit a bid for the good g . We assume that the bids of a buyer with valuation v is a number in the interval $[0, v]$.² The good will be allocated to the bidder who submitted the highest bid (with a lottery to determine the winner in a case of a tie).

The auction setup can be defined using a Bayesian game.³ In this game the players are the potential bidders, and the payoff of a player with valuation v is $v - p$ if he wins the good and pays p , and 0 if he does not get the good. The equilibrium concept can be also extended to the context of Bayesian games. In particular, in equilibrium of the above

²In general, buyers may submit bids that are higher than their valuations, but these strategies are dominated by other strategies, and their existence will not effect the equilibrium discussed in this paper.

³A formal definition and exposition of Bayesian games in our context will be presented in the full paper

game the bid of a player with valuation v is $(1 - \frac{1}{n})v$. Following the revelation principle, discussed in the economic mechanism design literature, one can replace the above-mentioned first-price auction with the following auction: each bidder will be asked to reveal his valuation, and the good will be sold to the bidder who reported the highest valuation; if agent i who reported valuation v' will turn out to be the winner then he will be asked to pay $(1 - \frac{1}{n})v'$. It turns out that reporting the true valuation is an equilibrium of that auction, and that it will yield (in equilibrium) the same allocation and payments as the original auction. It is convenient to consider the above *revelation mechanism*, since when facing any number of participants, a bidder's strategy in equilibrium will always be the same.

A first-price auction setup will be identified with a sequence of (Bayesian) games (g_1, g_2, \dots) where g_j is the Bayesian game associated with (the revelation mechanism of) first-price auction with $j + 1$ potential buyers. The definition of C -competitive strategies can now be applied to the above context as well.

The proof of the following appears in the full paper:

Theorem 4 *There exists an e -competitive strategy for the first-price auction setup.*

Our result can be obtained also if we consider standard first-price auctions, rather than the revelation mechanisms associated with them; nevertheless, this will require to allow a player to choose its action knowing the number of potential bidders.

Discussion

Some previous work in AI has attempted to show the potential power of decision-theoretic approaches that do not rely on classical game-theoretic analysis. In particular, work in theoretical computer science on competitive analysis has been extended to deal with rationality constraints (Tennenholtz 2001), in order to become applicable to multi-agent systems. We introduced competitive safety analysis, bridging the gap between the normative AI/CS approach and classical equilibrium analysis. We have shown that an observation, which is of great interest from a descriptive perspective to economists, can be extended and generalized to provide a powerful normative tool for computer scientists and AI researchers interested in protocols for non-cooperative environments. We have illustrated the use and power of competitive safety analysis in various contexts. We have shown general results about 2×2 games, as well as about games with many participants, and introduced the use of competitive safety analysis in the context of decentralized load balancing, leader election, and auctions. Notice that our work is concerned with a normative approach to decision making in multi-agent systems. We make no claims as for the applicability of this approach for descriptive purposes. Although there exists much literature on the failure of Nash equilibrium, it is still the most powerful concept for action prediction in multi-agent systems.

References

- Aumann, R. 1985. On the non-transferable utility value: A comment on the Roth-Shaper examples. *Econometrica* 53(3):667–677.
- Borodin, A., and El-Yaniv, R. 1998. *On-Line Computation and Competitive Analysis*. Cambridge University Press.
- Koutsoupias, E., and Papadimitriou, C. 1999. Worst-Case Equilibria. In *STACS*.
- Kraus, S. 1997. Negotiation and cooperation in multi-agent environments. *Artificial Intelligence* 94:79–97.
- Nisan, N., and Ronen, A. 1999. Algorithmic mechanism design. Proceedings of STOC-99.
- Rosenschein, J. S., and Zlotkin, G. 1994. *Rules of Encounter*. MIT Press.
- Roth, A. E. 1980. Values for games without side payments: Some difficulties with current concepts. *Econometrica* 48(2):457–465.
- Sandholm, T. W., and Lesser, V. 1995. Equilibrium Analysis of the Possibilities of Unenforced Exchange in Multiagent Systems. In *Proc. 14th International Joint Conference on Artificial Intelligence*, 694–701.
- Tennenholtz, M. 2001. Rational Competitive Analysis. In *Proc. of the 17th International Joint Conference on Artificial Intelligence*, 1067–1072.