Generalized correlated equilibrium for two-person games in extensive form with perfect information

Ferenc Forgó

Abstract A correlation scheme (leading to a special equilibrium called "soft" correlated equilibrium) is applied for two-person finite games in extensive form with perfect information. Randomization by an umpire takes place over the leaves of the game tree. At every decision point players have the choice either to follow the recommendation of the umpire blindly or freely choose some other action except the one suggested. This scheme can lead to Pareto-improved outcomes of some other correlated equilibria. Computational issues of maximizing a linear function over the set of soft correlated equilibria are considered and a linear-time algorithm in terms of the number of edges in the game tree is given for a special procedure called "subgame perfect optimization".

Keywords Correlation, Nash equilibrium, behavioral strategies, protocol

JEL Classification Number C72

F. Forgó
Department of Operations Research, Corvinus University of Budapest
Fővám tér 8, 1093 Budapest, Hungary
E-mail: ferenc.forgo@uni-corvinus.hu

1 Introduction

Aumann (1974) introduced the concept of correlated equilibrium (CE) for normal form finite games in order to Pareto-improve the payoffs a Nash (1951) equilibrium (NE) can provide. Later Aumann himself showed that CE perfectly conforms with Bayesian rationality when applied in a game theoretic context, Aumann (1987). The basic scenario CE is interpreted in, goes like this: an umpire does a lottery over the action (strategy) profiles according to a commonly known probability distribution. Then he suggests each player in secret what strategy to play. Suggestions are undisclosed to the rest of the players. A probability distribution over the strategy profiles is a CE if, in expectation, no player is motivated to play any other strategy but the one suggested. CE has proved to be very useful in several ways, in particular in realizing payoffs strictly better for all players than those of any NE's. Yet, for some games, CE cannot produce Pareto-better outcomes than NE.

Moulin and Vial (1978) are the first to depart from Aumann's protocol of a CE. They demand more commitment from each player by requiring that she pledge to follow the suggestion of the umpire blindly. If she does not want to use the services of the umpire, then she is free to choose from the available actions including the one the umpire would have suggested to her had she committed herself to take the advice of the umpire without any reservation. It is convenient to visualize the situation in the following way. The umpire puts his recommendation in a sealed red envelope and all the available actions in a sealed white envelope. The player has to decide: either choose the red envelope and do whatever is written inside, or choose the white envelope and take one of the actions in it. The probability distribution by which the umpire has selected the actions landing in the red envelopes is said to be a weak CE (WCE), if it is in no player's interest to choose the white envelope given that the other players pick the red one. As shown by Moulin and Vial (1978), for two-person games with at least three strategies per player, strictly better payoffs may be realized in a WCE than in any CE.

There are games, however, where even this setup does not help in achieving desirable outcomes. Forgó et al. (2005) introduced a new scenario in the context of extensive-form finite games with perfect information which can improve on NE outcomes in certain climate change negotiation games even in the case when CE and WCE cannot. The protocol is the following, using the "envelope" terminology. The umpire selects an action profile according to a commonly known probability distribution. He puts his recommendation in a red envelope for every player and all the other (not selected) actions in a white envelope. Each player then has to choose, not knowing what is inside, between the red envelope and the white one. If she picks the red one, she must do what is recommended inside. If she picks the white envelope, then she is free to choose any action included in it. The probability distribution for which, in expectation, it is in no player's interest to choose the white envelope, is called soft correlated equilibrium (SCE).

The above protocol can be applied to games in normal form in a straightforward way and it is the subject of another paper. This paper is concerned with the adaptation of the concept of CE to a special class of games in extensive form. Our main motivation is to create realizable protocols that are computationally treatable and lead to improvements of NE outcomes in equilibrium. Adjusting the idea of CE to extensive form games has been studied in various settings by Forges (1986, 1993), Myerson (1986), von Stengel and Forges (2007). Closest to the idea of SCE in extensive form games is the agent-form correlated equilibrium (AFCE) as defined by von Stengel and Forges (2007) which is in fact an adaptation of WCE for extensive games. The basic difference between our approach and that of von Stengel and Forges (2007) is that randomization takes place over the leaves of the game tree and not on moves in information sets. This can only be done because the analysis is restricted to games with perfect information where probabilities at the (singleton) information sets can uniquely be computed from terminal probabilities. Another avenue opens up, however, since subgame perfection can meaningfully be required. For extensive games with perfect information AFCE reduces to NE when correlation, stability as expressed by incentive constraints and subgame perfectness are required at the same time.

Since the set of SCE's is defined by linear inequalities (incentive constraints) the umpire can maximize a utility function over this set, thereby pursuing a common goal as expressed by the utility function in a self-enforcing way. This idea is discussed in Myerson (1985) in general and in Forgó et al. (2005) as applied in a model of climate change negotiations. To solve the optimization problems thus obtained in a computationally efficient way is also of concern. This issue is addressed at the end of the paper and low-cost algorithms are suggested for special cases. Computational aspects of correlated equilibria in more general settings have been extensively studied (see e.g. Gilboa and Zemel 1989, Urbano and Vila 2002, Papadimitriou and Roughgarden 2007, von Stengel and Forges 2007).

The paper is organized as follows: In Section 2, the conceptual foundations of different forms of correlated equilibria are discussed. Section 3 is devoted to combining correlated equilibria and subgame perfection. Section 4 addresses the issue of optimization over correlated equilibria. Section 5 concludes.

2 Correlated equilibria for extensive form games

Equilibria of games in extensive form are defined generally in terms of the normal form derived from the extensive form. Since CE and its variants have been defined for normal form games, in order to write down the incentive constraints of CE one needs to do this cumbersome conversion mostly leading to an exponential growth in problem size, the loss of the natural setting and interpretational ease the game has originally had.

Among the various proposals designed to overcome this difficulty the most appealing is the concept of extensive form correlated equilibrium (EFCE) as introduced by von Stengel and Forges (2007). EFCE is defined for extensive form games with perfect recall. In line with the original idea CE represents for normal form games, in the pre-play phase the umpire does a lottery over the moves at every information set according to a commonly known probability distribution and puts his suggestion in a sealed envelope. The recommendation belonging to a particular information set is handed over to the player whose turn is to move when she actually gets to this information set and she either takes the advice or makes a move of her own choice. Thus, most envelopes will never get opened. The set of probability distributions over moves at information sets is an EFCE if it is no player's interest to deviate from the recommendation given that the other players always follow the suggestion of the umpire. So, when a player considers deviation from the suggested move, she may choose any moves at her subsequent information sets.

Randomization can, however, take place in other ways as well. An example is AFCE, the agent-form correlated equilibrium where an "agent" is assigned to every information set and does the lottery only when the play gets to that particular information set and then hands over the envelope with the recommended move to the player whose turn is to move. The set of probability distributions is an AFCE if the recommended moves at any information set I maximize the expected payoff when recommendations at all other information sets are accepted and the suggested moves executed including those of the player who is to move at I. AFCE can rightly be thought of as the adaptation of WCE for extensive games.

In finite normal form games randomization over action(strategy) profiles or payoffs is generally the same thing. In extensive form games, payoffs occur at terminal nodes whose number is far less than that of strategies even if we confine ourselves to reduced normal forms and games with perfect information. In what follows we will consider the most simple case: finite two-person games with perfect information, no chance moves are allowed and all payoffs are assumed to be different.

If the game is of perfect information, then from probabilities assigned to terminal nodes a unique (conditional) probability distribution can be derived for each decision point. We will call these probability distributions behavioral probabilities, pretty much in line with the usual definition of behavioral strategies. On the other hand, if randomization takes place independently at each decision node (as in the case of AFCE), behavioral probabilities uniquely determine a probability distribution on the leaves of the game tree.

Formally, given a game tree $T = (V, E)$ with node set V and edge set E, a probability distribution p over the set of leaves L is said to be a tree-correlated strategy. Using this probability distribution known to every player, the umpire performs a lottery at each decision point and makes a recommendation for the player who is to move at that particular node of the game tree. Depending on the degree of commitment of the players to follow the umpire's suggestions, we will get the "behavioral" versions of CE, WCE and SCE which we will call tree-correlated, weak tree-correlated and soft tree-correlated equilibria and abbreviate them as TCE, WTCE and STCE, respectively. We also have to make assumptions about the reaction of the umpire to rejection of the suggested move. Allowing theoretically other possibilities as well, basically we will assume, if not stated otherwise, that he immediately withdraws from the game and lets the game proceed as if no correlation were permitted.

The randomization at a particular node takes place as follows. At every node, the umpire selects a move randomly according to the given (conditional) probabilities. Then we will consider three cases.

a) TCE: He tells the selected move as his recommendation to the player assigned to the node. The player either accepts the recommendation and acts accordingly, or does not accept it and is free to choose any moves. Then the game either ends by reaching a leaf or proceeds from the ensuing decision node with no umpire to suggest any moves thereafter.

b) WTCE: The umpire asks the player whether she wants to follow his suggestion blindly, or not. If she says yes, then tells her the suggested move which she will make. If not, then he lets her do whatever she wants to and withdraws from the game.

c) STCE: The same as for WTCE except that if the player does not follow the suggestion blindly, then the umpire tells her his suggested move which she must not choose, otherwise she can do anything. Again, in the latter case, the umpire withdraws.

For all scenarios (TCE, WTCE, STCE) a probability distribution over the leaves of the game tree is in equilibrium if no player can improve her expected payoff by doing anything else but follow the suggestion any time it is her turn to move.

Note that from that decision node on where the suggestion of the umpire was turned down a new game of perfect information begins whose outcome is assumed to be the (generically unique) subgame perfect Nash equilibrium obtainable by backward induction.

3 Subgame perfect soft tree-correlated equilibria

Subgame perfection is the most powerful refinement of equilibrium in extensive games. It makes sense if it is required of the three kinds of tree-correlated equilibria, i.e. as play proceeds probabilities representing a TCE, WTCE and STCE, respectively, when restricted to a subtree that may occur with positive probability and redefined as probabilities conditional on having reached the root of the subtree also form a TCE, WTCE and STCE, respectively, in the subgame. Subgame perfection prevents the umpire from losing credibility by modifying at some subsequent decision point the probabilities already made public at the outset. As it turns out, most of the time this is too much to demand, the generalization loses its power and nothing remains but Nash equilibria.

Proposition 1 WTCE (and therefore TCE as well) cannot produce Pareto-better subgame perfect outcomes than NE.

Proof Because we have assumed that all payoffs are different and there are no chance moves, backward induction gives a unique subgame perfect NE. Assume that the root is at the top of the game tree. Let us proceed from the bottom up as in backward induction. As long as no moves (edges) are excluded from consideration, which is the case for WCTE, the player making the last decision cannot get more by any randomization, independent of the degree of her commitment, than the maximum of her payoffs at the terminal nodes. Thus she gets the NE payoff. Then by going one level up in the tree, we can argue similarly and by backward induction this holds for the whole tree.□

Notice that if the requirement of subgame perfection is lifted, Proposition 2 fails to hold any more.

STCE, however, may give Pareto-better outcomes, as it will be shown later by an example.

According to the protocol of STCE, rejection can take place only once, when the umpire leaves the scene and subgame perfect NE will be realized via the application of backward induction. Given a game of perfect information Γ with game-tree $T = (V, E)$, node set V and edge set E, let a probability distribution p over the set of leaves L be a tree-correlated strategy. Since rejection can occur at every decision point, the same type of incentive constraints are associated with every node in the game tree. Assume that rejection first takes place at node N where it is player A's turn to move. Denote the subgame beginning at node N by G. The edges emanating from N are $e_1, ..., e_k$ leading to subgames $G_1, ..., G_k$. In subgame G_i the payoffs of player A at the leaves are denoted by $c_{i_1}, ..., c_{i_{r_i}}$, $i = 1, ..., k$. The terminal probabilities $p_{i_1}, ..., p_{i_{r_i}}$, $i = 1, ..., k$

are determined by the probability distribution p defined over the leaves L. So, subgame G_i is reached with probability $q_i = \sum_{s_i=1}^{s_i=r_i} p_{i_{s_i}}$, $i = 1, ..., k$ and game G is reached with probability $q = \sum_{i=1}^{i=k} q_i = \sum_{i=1}^{i=k} \sum_{s_i=1}^{s_i=r_i} p_{i_{s_i}} > 0$. The (conditional) probabilities by which play proceeds along the edges $e_1, ..., e_k$ are $\frac{q_1}{q}, ..., \frac{q_k}{q}$. The NE payoffs (of A) in the subgames $G_1, ..., G_k$ are denoted by $f_1, ..., f_k$. The expected payoff C_i of player A in subgame G_i if she obeys the suggestion of the umpire is

$$C_i = \frac{1}{q_i} \sum_{s_i=1}^{s_i=r_i} p_{i_{s_i}} c_{i_{s_i}} ,$$

and her expected payoff C in game G is

$$C = \sum_{i=1}^{i=k} \frac{q_i}{q} C_i = \frac{1}{q} \sum_{i=1}^{i=k} \sum_{s_i=1}^{s_i=r_i} p_{i_{s_i}} \; c_{i_{s_i}} .$$

If she does not want to obey blindly, then in the subgame of her choice, she will get, by assumption, the NE payoff. Therefore, her expected payoff will be

$$\sum_{i=1}^{i=k} \frac{q_i}{q} \max_{j \neq i} f_j .$$

Thus, after multiplying both sides by q, we will get the incentive constraint at node N, a linear inequality

$$\sum_{i=1}^{i=k} \sum_{s_i=1}^{s_i=r_i} p_{i_{s_i}} \; c_{i_{s_i}} \geq \sum_{i=1}^{i=k} \sum_{s_i=1}^{s_i=r_i} p_{i_{s_i}} \; \max_{j \neq i} f_j \; . \tag{1}$$

When it is player B's turn to move we will get the same type of incentive constraints.

Example 1 Take the centipede game (see Osborne and Rubinstein (1996)) of length 4 (the players make decisions 4 times, taking turns and may choose either "stop" or "continue"). Player A starts, payoffs when she stops are $(1, 0)$ and $(3, 2)$, When player B stops, payoffs are $(0, 3)$ and $(2, 5)$. Player B makes the last decision and if she opts to continue, then the game ends with the final payoff $(5, 4)$. Subgame perfect STCE's are feasible solutions of the following system of inequalities in nonnegative variables:

Table 1

Defining inequalities of STCE of a four-leg centipede game

p_1	p_2	p_3	p_4	p_5		
			-1	1	$\leq$	$0,$
		-1	1	-2	$\leq$	$0,$
	-1	1	-2	-1	$\leq$	$0,$
-1	1	-2	-1	-4	$\leq$	$0,$
1	1	1	1	1	$=$	$1.$

Terminal nodes of the game-tree are reached with probabilities $p_1, ..., p_5$ and to each decision point there belongs an incentive inequality. For example, if player A is to move for the second time, then in the subgame her expected payoff she gets when obeying the suggestion of the umpire blindly is

$$\frac{p_3}{p_3+p_4+p_5}3+\ \frac{p_4}{p_3+p_4+p_5}2+\frac{p_5}{p_3+p_4+p_5}5,$$

whereas if she disobeys and the umpire withdraws from giving advice, then she gets the NE payoffs 3 with probability $\frac{p_4+p_5}{p_3+p_4+p_5}$ and 2 with probability $\frac{p_3}{p_3+p_4+p_5}$.This is how we get the incentive constraint

$$3p_3 \quad +2p_4 \quad +5p_5 \ \geq \ 2p_3 \quad +3p_4 \quad +3p_5$$

which is the second inequality in Table 1. The other inequalities are derived likewise.

It is easy to see that $p_4 = \ 1/2$, $p_5 = \ 1/2$ all other p_j's are zero, satisfies the incentive constraints and is thus a STCE. It is also easy to see that no matter how long the centipede game is, suggesting to stop or continue at the last decision point with probabilities $\frac{1}{2}$ is always a STCE.

Assuming that rejection of blindly following the umpire's suggestion at a decision point makes him completely withdraw and let the game proceed unattended towards NE payoffs may seem too restrictive. In fact, the incentive constraints can be generalized if we only assume that there is a function h_N at every decision node N defined over the set of all finite games of perfect information with root N, assigning an outcome (payoff) for each player if rejection takes place at N. We might as well call this function as a "threat" function since this is what every player has to weigh against blind obedience when choosing the white envelope. A pessimistic approach for player A (or B) is to take the *maxmin* value from her point of view of the subgame beginning at N. Since this is never more than her NE payoff, the incentive constraints will not be more restrictive (possibly less) than the ones defined by (1) leaving more room for achieving desirable outcomes. Being "optimistic", i.e. counting on better than NE outcomes can make the set of STCE's empty since in this case some players may rightly think that they can be better off by rejecting blind obedience.

We will have a different scenario if protocols of AFCE and STCE are combined in the following way. Deviation from the suggested move can only take place once (AFCE) and the deviating player can only choose from moves other than the one suggested (STCE). This scenario models a situation where the umpire does not get offended by the rejection of his recommendation but keeps on giving advice in order to promote better outcomes. The restriction that deviation can occur only once can be interpreted as some sort of stability: In equilibrium (we call it soft agent-form correlated equilibrium, SAFCE for short) it is not worth (in expectation) for a player to dismiss the advice of her agent if all other agents's recommendations are accepted.

Incentive constraints of SAFCE can also be obtained at each decision node by comparing the expected payoff of full obedience with the best possible expected payoff if deviation occurs. We can proceed as we did in the case of STCE keeping the notation as well. Assume that rejection takes place at node N where it is player A's turn to move. If at N, edge e_i is recommended and player A does not accept blind obedience, then the largest payoff player A can get in subgame G_i is

$$\sum_{r=1}^{r=k} \frac{q_r}{q} \max_{j \neq i} C_j,$$

and thus the incentive constraint is

$$\sum_{r=1}^{r=k} \frac{q_r}{q} C_r \geq \sum_{r=1}^{r=k} \frac{q_r}{q} \max_{j \neq i} C_j,$$

which becomes after multiplying both sides by q

$$\sum_{r=1}^{r=k} q_r C_r \geq \sum_{r=1}^{r=k} q_r \max_{j \neq i} C_j \ .$$

This is, however, not a linear inequality, a major obstacle in the way of practical computations. It is easy to see that the unique NE gives rise to a SAFCE, thus local search starting from the NE may lead to other SAFCE's. In fact, the set of SAFCE's is always larger than that of the NE's.

Proposition 2 There are other SAFCE's than the unique NE.

Proof Assume that the game-tree is of length L and we are at a decision point D one move away from termination (root of a subtree of length 1) where player A is to select one from the leaves $l_1, ..., l_k$ with payoffs $b_1, ..., b_k$ for player A. The incentive constraint at D is now

$$\sum_{i=1}^{i=k} p_i b_i \geq \sum_{i=1}^{i=k} p_i \max_{j \neq i} b_j$$

where $p_1, ..., p_k$ are the probabilities of moves. Without loss of generality we may assume that $b_1 > b_2 > ... > b_k$. It is easy to verify that any probability distribution with $p_1 \geq \frac{1}{2}$ satisfies the incentive constraint. Choose and fix one such probability distribution with $\frac{1}{2} \leq p_1 < 1$ for every subtree of length 1. Define now a game-tree of length $L - 1$ where each leaf represents a subgame of length 1 with payoffs that are the players's expected payoffs using the fixed probabilities in that particular subtree. We can now repeat the procedure until we reach the root of the game-tree. What we will get finally is a SAFCE, terminal probabilities are products of probabilities along the paths leading from the root to leaves. Since the first term of this product is less than 1, it cannot be that of the unique NE for which every probability along its path is 1.□

We saw in the proof that the SAFCE constructed by backward induction satisfies the incentive constraint but there is no guarantee that any SAFCE satisfying the incentive constraints given in terms of terminal probabilities can be obtained in such a way.

Example 2 For illustration, consider the centipede game in Example 1. Denote the first and second decision points of players A and B by $A1, A2$ and $B1, B2$, respectively, and the probabilities of "continue" at $A1, B1, A2, B2$ by p, q, r, s. The incentive constraints at the decision points in reverse order are as follows:

$$\begin{aligned}
B2 &: \quad 4s + 5(1-s) \geq 5s + 4(1-s), \\
A2 &: \quad 3(1-r) + 2r(1-s) + 5rs \geq 3r + (5s + 2(1-s))(1-r), \\
B1 &: \quad 3(1-q) + 2(1-r)q + 5r(1-s)q + 4rsq \geq \\
&\quad 3q + (2(1-r) + 5r(1-s) + 4rs))(1-q), \\
A1 &: \quad 1 - p + 3q(1-r)p + 2qr(1-s)p + 5qrsp \geq \\
&\quad p + (3q(1-r) + 2qr(1-s) + 5qrs)(1-p).
\end{aligned}$$

These inequalities are all nonlinear except for the first one. This remains so even if we express p, q, r, s in terms of terminal probabilities p_1, p_2 , p_3, p_4, p_5. In particular

$$s = \frac{p_5}{p_4 + p_5},$$

$$r = \frac{p_4 + p_5}{p_3 + p_4 + p_5},$$

$$q = \frac{p_3 + p_4 + p_5}{p_2 + p_3 + p_4 + p_5},$$

$$p = p_2 + p_3 + p_4 + p_5.$$

In order to find a particular SAFCE, set $s = 1/4$ which satisfies inequality $B2$. Then the expected payoffs of the players at $B2$ are $(11/4, 19/4)$. With these numbers fixed, the incentive constraint at $A2$ is

$$3(1-r) + \frac{11}{4}r \geq 3r + \frac{11}{4}(1-r).$$

Take now (arbitrarily) $r = 1/3$ satisfying the above inequality and compute the expected payoffs at $A2$ which turn out to be $(35/12, 35/12)$. The incentive constraint at $B1$ is now

$$3(1-q)+\frac{35}{12}q \geq 3q+\frac{35}{12}\ (1-\ q)$$

Set $q=1/2$ satisfying the above inequality and compute the expected payoffs at $B1 : (35/24, 71/24)$. The incentive constraint at $A1$ is

$$1-p+\frac{35}{24}p \geq p+\frac{35}{24}\ (1-\ p)$$

Pick a p satisfying this inequality, say $p=3/4$ and then we have the behavioral probabilities $p=3/4$,$q=1/2$,$r=1/3, s=1/4$ and terminal probabilities computed from them $p_1=8/32, p_2=12/32, p_3=8/32, p_4=3/32, p_5=1/32$, a SAFCE. This happens to be a STCE but generally it does not have to be so.

A careful look at the definition and the incentive constraints of STCE reveals that it is important that every subgame have a unique NE payoff for both players. Exclusion of chance moves and identical terminal payoffs are just sufficient conditions to ensure the uniqueness of NE payoffs. In comparison, this condition can be abandoned for SAFCE since NE payoffs in subgames play no role in the construction of incentive constraints.

4 Optimization over the set of tree-correlated equilibria

Usually there are a lot of (generalized) correlated equilibria satisfying the incentive constraints. If we want to combine self-enforcing stability as expressed by incentive constraints derived from various protocols, and optimality in the form of a linear objective function whose maximum over the set of correlated equilibria represents a desirable outcome from the umpire's perspective, then we have an optimization problem on our hands. In this case, in addition to the two payoffs, there is one more number assigned to every leaf $l \in L$, the utility of the umpire if play terminates at l. The umpire's utilities are also common knowledge. Since the STCE's are defined by linear inequalities, the optimization problem is a linear programming problem.

If we want to optimize over the set of SAFCE's, then the nonlinear constraints may cause severe problems. Optimization becomes much easier, however, if we do "subgame perfect optimization". The idea is based on backward induction and we will proceed in the same way we did when we determined a particular SAFCE except that we do not choose arbitrary feasible probabilities but the probabilities will be obtained via optimization over the set of probabilities satisfying the incentive constraint.

Here is a rather informal description of how subgame perfect optimization works. Assume that we are at a decision point D, one move away from termination. If the umpire wants to pick a SAFCE maximizing a linear function, then at this stage, he can only decide over the probabilities of going from D to leaves in this simple subtree. There is only one incentive inequality and the objective function is linear in the probabilities allowing to compute the optimal solution efficiently. Once this problem is solved, we will think of D as a leaf of the game representing the ensuing subgame together with the optimal payoffs that have been obtained in the subgame beginning with D. We can proceed in this fashion from the bottom up until we reach the root of the tree. Since the set of SAFCE's in each subgame is never empty (see Proposition 2), we will finally have a subgame perfect optimal solution. Of course the objective function value of this solution cannot be better (usually it is worse) than the one obtained by optimizing over all SAFCE's. Modification of probabilities as play proceeds from the root towards the leaves undermines the credibility of the umpire since probabilities (the tree correlated strategy) are announced at the outset to enable players make their decisions based on expected gains and losses.

Example 3 For illustration take the four-leg centipede game of Example 3 and 4 again and the linear objective function given in terms of terminal probabilities:

$$p_1 + 3p_2 + 5p_3 + 7p_4 \to \max$$

To compute a subgame perfect optimal solution we first convert the terminal probabilities to products of behavioral probabilities and get the objective function in the following form

$$1 - p + 3p(1-q) + 5pq(1-r) + 7pqr(1-s).$$

To begin with, solve the umpire's optimization problem at the last decision point $B2$

$$\begin{aligned} 7(1-s) &\to \max \\ 4s + 5(1-s) &\geq 5s + 4(1-s), \\ 0 &\leq s \leq 1, \end{aligned}$$

where s denotes the probability of continuing. The optimal solution is $s = 0$. Thus, by backward induction, everybody knows that if play gets to $B2$, then the payoffs for the players will be 2 and 5, and the objective function value will be 7. This is so because deviation can only take place only once by the definition of SAFCE. Therefore at $A2$ the umpire has to solve the problem

$$\begin{aligned} 7r + 5(1-r) &\to \max \\ 2r + 3(1-r) &\geq 3r + 2(1-r), \\ 0 &\leq r \leq 1 \end{aligned}$$

whose solution is $r = 1/2$.. Thus everybody at the decision point $B1$ thinks that if player B plays "continue", then the expected payoffs of the players will be 5/2 and 7/2 ,respectively and the umpire's objective function value will be 6. Stopping, as is given in the description of the game, gives payoffs 0 and 3, and objective function value 3. So the umpire solves the problem

$$\begin{aligned} 6q+3(1-q) &\rightarrow \max \\ 3.5q+3(1-q) &\geq 3q+3.5(1-q), \\ 0 &\leq q \leq 1. \end{aligned}$$

The solution is $q = 1$. Similar calculation shows that $p = 1$, that is, player A at her first decision point $A1$ is advised to play "continue". Putting everything together, it turns out that in the SAFCE we have just determined both player A and player B are first recommended to play "continue", then player A is advised to play "stop" or "continue" with probability 1/2, and finally it is suggested for player B to stop if the play ever gets there. The subgame perfect optimal behavioral probabilities are

$$p = 1, q = 1, r = \frac{1}{2}, s = 0$$

or, in terminal probability form

$$p_1 = 0, p_2 = 0, p_3 = \frac{1}{2}, p_4 = \frac{1}{2}, p_5 = 0.$$

The objective function value is 6. This happens to be optimal over the set of SAFCE's but as mentioned earlier, overall optimality cannot be guaranteed.

Proposition 3 Subgame perfect optimization over the set of SAFCE's with a linear objective function finds an optimal solution in time linear in the number of edges of the game tree.

Proof When doing subgame perfect maximization of a linear objective function over the set of SAFCE's, at every decision point one needs to solve a linear programming problem with two constraints: the incentive constraint at that particular node and the normalizing equality (probabilities must add up to 1). As Dyer (1984) has shown, this problem can be solved in time linear in the number of variables, that is, edges emanating from the node. Thus the subgame perfect optimization over the whole tree can be done in time $\sum_{i=1}^{d} n_i$ where n_i denotes the number of edges branching off at node i and d is the number of decision points.□

It is worth noting that the objective function, at every decision point, has just as many variables as the number of edges going out (one variable can be eliminated right away using the normalizing equation). Therefore for small number of variables (as in the case of not too long centipede games) practically any nonpathologic objective function is computationally treatable just as in the case of several variables with an objective function that is concave in the variables representing the behavioral probabilities at the particular decision point.

5 Concluding remarks

We have demonstrated that the protocol of a generalization (soft correlated equilibrium) can meaningfully be adapted to games in extensive form with perfect information. Linear utility functions can effectively be maximized over the set of equilibria.

Several questions have remained unanswered and quite a few areas unexplored. The main challenge is to characterize the set of games whose correlated equilibria outcomes cannot be improved by soft correlated equilibria. Formulation of soft correlated equilibrium when the original game is of imperfect and/or incomplete information also calls for further research. Though extension of the ideas to more than two players seems quite straightforward at certain places, caution must be exercised and careful study should decide where and how the results of the two-person case can be carried over to the many-player situation.

Acknowledgements

Special thanks are due to János Fülöp whose help during the preparation of this paper was most valuable. I also thank Sándor Bozóky for helping in solving some nonconvex numerical examples. However, I am solely responsible for any mistakes and incorrect statements remaining. Research was sponsored by grant OTKA 72856.

References

Aumann RJ (1974) Subjectivity and correlation in randomized strategies. Journal of Mathematical Economics 1:67-96

Aumann RJ (1987) Correlated equilibrium as an expression of Bayesian rationality. Econometrica 55:1-18

Dyer ME (1984) An $O(n)$ algorithm for the multiple choice-knapsack linear program. Mathematical Programming 29:57-63

Forges F (1986) An approach to communication equilibria. Econometrica 54:323-358

Forges F (1993) Five legitimate definitions of correlated equilibrium in games with incomplete information. Theory and Decision 35:277-310

Forgó F, Fülöp J, Prill M (2005) Game theoretic models for climate change negotiations. European Journal of Operations Research 160:252-267

Fülöp J (2007) Private communication

Gilboa I, Zemel E (1989) Nash and correlated equilibria: some complexity considerations. Games and Economic Behavior 1:80-93

Moulin H,Vial J-P (1978) Strategically zero-sum games: the class of games whose completely mixed equilibria cannot be improved upon. International Journal of Game Theory 7:201-221

Myerson R (1985) Mechanism design by an informed principal. Econometrica 47:61-73

Myerson R (1986) Multistage games with communication. Econometrica 54:323-358

Nash JF (1951) Noncooperative games. Annals of Mathematics 54:286-295

Osborne MJ, Rubinstein A (1996) A course in game theory. The MIT Press, Cambridge MA

Papadimitriou CH, Roughgarden T (2007) Computing correlated equilibria in multi-player games. Available at http://theory.stanford.edu/~tim/papers/cor.pdf

Urbano A, Vila JE (2002) Computational complexity and communication: coordination in two-player games. Econometrica 70:1893-1927

von Stengel B, Forges F (2007) Extensive form correlated equilibrium: Definition and computational complexity. CDAM Research Report LSE-CDAM-2006-04