

Fast algorithms for finding proper strategies in game trees

Peter Bro Miltersen* and Troels Bjerre Sørensen*

July 6, 2007

Abstract

We show how to find a *normal form proper* equilibrium in behavior strategies of a given two-player zero-sum extensive form game with imperfect information but perfect recall. Our algorithm solves a finite sequence of linear programs and runs in polynomial time. For the case of a perfect information game, we show how to find a normal form proper equilibrium in linear time by a simple backwards induction procedure.

*Department of computer science, University of Aarhus, Denmark. Email {bromille,trold}@daimi.au.dk. Research supported by *Center for Algorithmic Game Theory*, funded by the Carlsberg Foundation.

1 Introduction

It is well known that Nash equilibria of *matrix games* (i.e., two-player zero-sum games in normal form) coincide with pairs of minimax and maximin mixed strategies and can be found efficiently using linear programming. However, in many realistic situations where it is desired to compute prescriptive strategies for games with hidden information, the game is given in *extensive form*, i.e., as a *game tree* with a partition of the nodes into *information sets*. Each information set describes a set of nodes mutually indistinguishable for the player to move. One may analyze an extensive form game by converting it into normal form and then analyzing the resulting matrix game. However, the conversion from extensive to normal form incurs an exponential blowup in the size of the representation. Koller, Megiddo and von Stengel [6] showed how to use *sequence form* representation to compactly represent and efficiently compute maximin behavior strategies for two-player extensive-form zero-sum games with imperfect information but perfect recall by solving linear programs of size linear in the size of the game trees, and avoiding the conversion to normal form. The method of Koller, Megiddo and von Stengel has been used for constructing prescriptive strategies for concrete, often very large games, to be used in game playing software. In particular, it was applied to solve various variants of two-player poker [1, 4] containing millions of information sets. The efficient algorithm, based on the sequence form representation and implemented using state of the art linear programming software was essential for obtaining solutions for games this large.

Despite its widespread use, the strategies computed by the algorithm suffer from a certain deficiency, first pointed out by Koller and Pfeffer [8]: While the strategy computed by the algorithm is a correct maximin strategy and thus guaranteed to attain an expected payoff of at least the game-theoretic value of the game considered against any counter strategy, it does *not* necessarily prescribe sensible play in any particular situation encountered during the game. Indeed, since the strategy computed is not attempting to achieve a payoff better than the value of the game, a player playing by the strategy will gladly give back any “gift” he receives from his opponent. The deficiency may be illustrated even with perfect information games. Consider the two perfect information games in Figure 1 (payoffs at leaves are paid by Player 2 to Player 1). The value of the game (a) is 0 and the strategy profile (i.e., a strategy for each player) choosing the actions indicated form a Nash equilibrium. However, if Player 1 gets to move, he should surely not perform the action U indicated, as the mistake of Player 2 has enabled him to achieve a payoff of 1 by choosing action D. Similarly, the value of the game (b) is 0 and the choices indicated form a Nash equilibrium. However, it does not seem sensible for Player 1 to open the game by choosing U as indicated. Rather than immediately achieving the payoff of 0, he should choose D, let Player 2 move and hope that Player 2 will make a mistake and choose L, resulting in a payoff of 1. When a perfect information game is solved using standard backwards induction (minimax evaluation), mistakes of type (a) are automatically avoided and it is straightforward




Figure 1: (a) Not subgame perfect. (b) Not admissible.

to break ties in the minimax evaluation so as to avoid mistakes of type (b). However, if the game is solved using the Koller-Megiddo-von Stengel algorithm, the "bad" equilibria may easily be given as output, and if the game has imperfect information, straightforward techniques for avoiding them do not apply.

Fortunately, in game theory, *refinements* of Nash equilibrium were defined with exactly the purpose of eliminating "insensible" Nash equilibria of extensive form games. For a comprehensive account, see the monograph of van Damme [17]. Using terminology from this theory, the equilibrium in (a) is not *subgame perfect*, since its restriction to the subgame starting with the position where Player 1 gets to move is not a Nash equilibrium. The equilibrium in (b) *is* subgame perfect, but not *admissible*, as the strategy it prescribes for Player 1 is weakly dominated by his other possible strategy. An attractive equilibrium refinement is *quasi-perfect equilibrium* which was introduced by van Damme [16]. A quasi-perfect equilibrium is guaranteed to be admissible as well as *sequential*, a non-trivial refinement of subgame perfection for imperfect information games due to Kreps and Wilson [9]. In previous work [13, 14], we showed how to modify the Koller-Megiddo-von Stengel linear programs so that a quasi-perfect equilibrium is computed. This eliminates the undesirable behavior in all the examples pointed out by Koller and Pfeffer.

However, while insisting on a quasi-perfect equilibrium thus eliminates most of the cases of "returning gifts" in the computed equilibria, there are still games that are not solved in a satisfactory way. Here, we show an example of an equilibrium of a fairly natural extensive-form game we call *Matching Pennies on Christmas Morning* and first presented in [12].

The game is as follows. In the standard Matching Pennies game, Player 2 (Bob) hides a penny and Player 1 (Alice) has to guess if it is heads or tails up. If she guesses correctly, she gets the penny. If played on Christmas morning, we add a *gift option*: After Player 2 has hidden his penny but before Player 1 guesses, Player 2 may choose to publicly give Player 1 a gift of one penny, in addition to the one Player 1 will get if she guesses correctly. The extensive form of this game as well as the pair of maximin/minimax behavioral strategies computed by the game theory software tool Gambit [11] using the Koller-Megiddo-von Stengel algorithm are given in Figure 2. We see that if Player 1 does not receive a gift, the strategy computed suggests that she randomizes her guess and guesses *heads* with probability $\frac{1}{2}$ and *tails* with probability $\frac{1}{2}$. This is indeed the strategy we expect to see. On the other hand if Player 1 *does* receive a gift, the strategy computed suggests that she guesses *heads* with probability 1. This does not seem sensible. Indeed, if she had randomized her guess, as in the "no-gift" scenario, her worst case conditional expected payoff, conditioned by the observed fact that she received the gift,

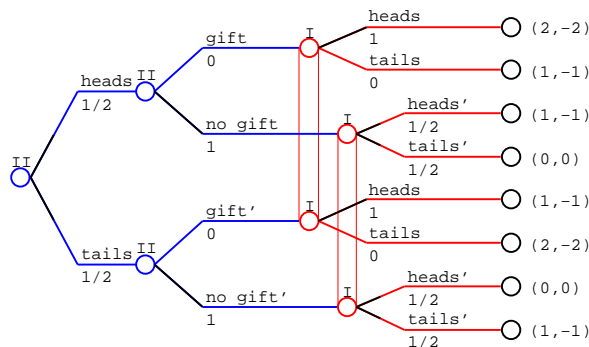


Figure 2: *Matching Pennies on Christmas Morning* - "bad" equilibrium (labels on arcs are behavior probabilities)

would be guaranteed to be at least a penny and a half. With the strategy computed, the worst case conditional expected payoff is only a penny. This will be the payoff in the case where the strategy of Player 2 happens to be the pure strategy of hiding the penny tails up and giving the gift. The “bad” equilibrium is quasi-perfect, and the algorithm of [13] and in fact all previously described variants of the Koller-Megiddo-von Stengel algorithm may give it as output.

The only standard equilibrium refinement we are aware of that only permits the “sensible” equilibrium of Matching Pennies on Christmas Morning is the notion of *proper equilibrium* of Myerson [15]. An equilibrium in mixed strategies for a bimatrix game is said to be proper if it is a limit point of a sequence of ϵ -proper completely mixed strategy profiles for $\epsilon \rightarrow 0+$. Here, a strategy profile is said to be completely mixed if it prescribes strictly positive probability to every pure strategy. It is said to be ϵ -proper if the following property is satisfied: If pure strategy x_i is a better reply than pure strategy x_j against the mixed strategy the profile prescribes to the other player, we have $p(x_j) \leq \epsilon p(x_i)$ where $p(x_k)$ is the probability prescribed to pure strategy x_k . An equilibrium in behavior strategies for an extensive form two-player game with perfect recall is defined to be *normal form proper* if it is behaviorally equivalent to a proper equilibrium of the corresponding normal form game. Using fixed point arguments, Myerson showed the existence of a normal form proper equilibrium for any such game. It is easy to see that a normal form proper equilibrium is also a Nash equilibrium.

The intention of the definition of proper equilibrium is to capture the behavior of players that assume that their opponents sometimes make mistakes, but with negligible probability and *in a rational manner*: Players assume that their opponents play more costly suboptimal strategies with significantly smaller probability than less costly suboptimal strategies, but that any suboptimal strategy is played with infinitesimally small probability. It can be seen that the unique normal form proper equilibrium of Matching Pennies on Christmas Morning is the one where Player 1 guesses uniformly at random, also after having received a gift. The intuitive reason is that even after observing that Player 2 has made a mistake by giving the gift, Player 1 assumes that Player 2 in all other respects behaves rationally. In particular, she assumes that he will exploit any bias in the guess she makes, and hence she has to avoid any such bias.

Kohlberg and Mertens [5] and van Damme [16, 17] established a number of attractive properties of normal form proper equilibrium, yielding additional motivation for computing it. First, any normal form proper equilibrium in behavior plans is also a quasi-perfect equilibrium in behavior plans [5, 16], so examples where the “sensible” solutions are guaranteed by quasi-perfection (see [13] for non-trivial such examples) are still solved in a satisfactory way when a normal form proper equilibrium is computed. Also, for the case of a zero-sum game, van Damme [17, Theorem 3.5.5] has shown that the set of proper equilibria is a Cartesian product $D \times D'$, where D is a polytope of mixed strategies for Player 1 and D' is a polytope of mixed strategies for Player 2. This nice fact makes it possible to extend the properness terminology from strategy profiles to strategies: We define a proper mixed *strategy* for Player 1 (resp., 2) to be an element of D (resp., D') and a proper behavior strategy to be a behavior strategy behaviorally equivalent to a proper mixed strategy. Note that this is completely analogous to the case of Nash equilibria which for the case of a zero-sum game is the Cartesian product of the maximin strategies of Player 1 and the minimax strategies of Player 2, by von Neuman’s min-max theorem.

The above discussion motivates the computation of proper strategies when prescriptive behavior in an extensive form zero-sum game is desired. In the present paper, we show how to extend the approach of Koller, Megiddo and von Stengel and do such a computation in polynomial time. First, we study the case of *imperfect information* games with perfect recall and show that the normal form proper equilibria of such games can be completely characterized by a procedure involving an iteratively defined sequence of linear programs derived from the linear

programs for maximin behavior plans in sequence form described by Koller, Megiddo and von Stengel. Each linear program in the sequence has a number of variables and constraints which is at most the size of the game tree and the number of programs in the sequence is also at most the size of the game tree. This establishes our first main theorem:

Theorem 1 *A normal form proper equilibrium in behavior plans for a given extensive form two-player zero-sum game of imperfect information but perfect recall can be found in polynomial time in the size of the game tree.*

The sequence of linear programs we construct is analogous to the sequence of linear programs arising in *Dresher's procedure* [3] established by van Damme [17] as characterizing the proper equilibria of a matrix game and our proof of correctness is based on similar ideas as van Damme's proof of this fact. Thus, our algorithm may be seen as a sequence form version of Dresher's procedure. However, the programs we devise are not in general equivalent to the linear programs arising in Dresher's procedure for the corresponding normal form game, and indeed, for some games our sequence of programs is exponentially shorter than the sequence of programs arising in Dresher's procedure.

Our main motivation for computing proper equilibria is computing sensible prescriptive behavior in imperfect information games. But, somewhat surprisingly, it turns out that normal form proper equilibrium is an interesting solution concept even for the case of *perfect* information games. As an example, consider the game given in Figure 3. The value of the game is 0 and Player 1 is guaranteed to obtain this value no matter what she does. However, if she chooses action U and her opponent makes a mistake, she will receive a payoff of 1. On the other hand, if she chooses action D and her opponent makes a mistake, she will receive a payoff of 2. In the unique normal form proper equilibrium for this game, Player 1 chooses U with probability $\frac{2}{3}$ and D with probability $\frac{1}{3}$. An intuitive justification for this behavior strategy is as follows. Player 1 must imagine being up against a Player 2 that cannot avoid sometimes making mistakes - otherwise her choice is irrelevant. On the other hand, she should assume that Player 2 is still a rational player who makes an effort to avoid making mistakes. In particular, she should assume that Player 2 makes a strategic decision about whether to train to avoid making the mistake in position h or to train to avoid making the mistake in position h' , prior to playing the game. Thus, the strategy of Player 1 should not be pure. In particular, if she chooses D with probability 1 (as is surely tempting), Player 2 may respond by concentrating all his efforts to avoid making mistakes in h' . Then, Player 1 will not get her "fair share" of payoff from Player 2's mistakes. This reasoning is somewhat analogous to the heuristic reasoning suggesting why it is a good idea for expert chess players to randomize which opening to select against other expert players, despite the fact that chess is a perfect information game. The above discussion motivates our second main result:

Theorem 2 *A normal form proper equilibrium in behavior strategies for a given extensive form two-player zero-sum game of perfect information can be found in linear time.*




Figure 3: *Up or Down?*

The procedure we exhibit is a *backwards induction* procedure, refining the standard backwards induction procedure (minimax evaluation) for computing a subgame perfect equilibrium. As a curious example, applying the procedure to *tic-tac-toe* one finds that in any normal form proper equilibrium of this game, the game is opened by selecting the middle square with probability $\frac{1}{13}$.

1.1 Related research

In his monograph, van Damme established that a procedure due to Dresher [3] characterizes the proper equilibria of a matrix game. One may apply this procedure to an extensive form game by first converting the game to normal form (incurring an exponential blowup in the size of the representation) and then applying the procedure. As described by van Damme (and Dresher), the procedure uses exponential time as it involves explicitly enumerating the vertices of a polytope defined by a set of inequalities, so this would yield a doubly exponential time procedure. In previous work [12], we showed how to relatively easily modify the procedure so that it becomes polynomial time, yielding a polynomial time procedure for computing a proper equilibrium of a *matrix* game and hence a singly exponential time procedure for computing a normal form proper equilibrium for a two-player zero-sum extensive form game with perfect recall. The procedure we present in this paper is much more subtle. It subsumes the procedure presented in [12] in the following sense: If a matrix game is represented as an extensive form game of depth two with one information set for each player and the procedure of the present paper is applied to this representation, we get essentially the same sequence of linear programs as when applying the procedure of [12] to the matrix game.

Computing a proper equilibrium of a *bimatrix* game remains an elusive problem (whose solution is a prerequisite for even considering computing a normal form proper equilibrium for a two-player extensive form non-zero sum game). Yamamoto [19] presents a numerical procedure for performing such a computation. It involves solving certain differential equations numerically and it is not clear under which circumstances it can be formally guaranteed to compute (or converge to) a proper equilibrium. One should note that since proper equilibrium refines Nash equilibrium, it is PPAD-hard to compute it, even for the bimatrix case, by the breakthrough result of Chen and Deng [2]. Given this, we find the following problem very interesting.

Open Problem 1 *Is computing a proper equilibrium of a bimatrix game PPAD-easy?*

To illustrate that the answer to this is far from obvious, we want to point out that no “pivoting”-style algorithm a la Lemke-Howson is known for computing a proper equilibrium of a bimatrix game. In fact, as far as we know, there is no known theorem guaranteeing a proper equilibrium using rational probabilities only of a given bimatrix game with integer payoffs. Lacking such a theorem, it is not even clear what “computing” a proper equilibrium means as one would have to settle for an approximation. But it is not clear to us if the standard relaxation of Nash equilibrium to ϵ -equilibrium has a meaningful counterpart for the case of proper equilibrium.

1.2 Organization of Paper

We present our procedure for imperfect information games in Section 2 and the procedure for perfect information games in Section 3. Due to lack of space, the proofs of correctness are in appendices. Appendix A presents the necessary background on extensive form games, the sequence form, and the result of Koller, Megiddo and von Stengel. Appendix B presents the

proof of correctness of the procedure for imperfect information games. Appendix C presents the proof of correctness of the procedure for perfect information games.

2 The procedure for imperfect information games

Let G be a two-player zero-sum extensive form game with perfect recall played between Player 1, trying to maximize payoff and Player 2, trying to minimize payoff. Our construction uses the *sequence form* of G and is based on the linear programming characterization of Nash equilibria in realization plans due to Koller, Megiddo and von Stengel [6]. A summary of this characterization is given in Appendix A. In particular, G is given by a payoff matrix A , and realization plan constraint matrices E and F for Player 1 and Player 2 respectively. We define a series of linear programs where the coefficients of each linear program depend on the solutions of previous linear programs. We group the linear programs in pairs, denoting a pair as a round, the first being round 0. In round 0, we first consider the following linear program $P^{(0)}$ which is the linear program devised by Koller, Megiddo and von Stengel for characterizing the maximin realization plans for Player 1 and computing the value of the game.

$$\begin{aligned} & \underline{P^{(0)}} : \\ \max_{x,q} & \quad f^\top q \\ \text{s.t.} & \quad -A^\top x + F^\top q \leq 0 \\ & \quad Ex = e \\ & \quad x \geq \mathbf{0} \end{aligned} \tag{1}$$

The vector variable x is indexed by action sequences of Player 1 and in the optimum solution describes a maximin realization plan for Player 1. The vector q is indexed by information sets of Player 2. We let $v^{(0)}$ be the value of the optimal solution. This is the value of the game, by the result of Koller, Megiddo and von Stengel. Next, we consider the following program $Q^{(0)}$.

$$\begin{aligned} & \underline{Q^{(0)}} : \\ \max_{x,q,u,s} & \quad \mathbf{1}^\top u \\ \text{s.t.} & \quad -A^\top x + F^\top q + u \leq \mathbf{0} \\ & \quad Ex - es = \mathbf{0} \\ & \quad f^\top q - v^{(0)}s = 0 \\ & \quad \mathbf{0} \leq u \leq \mathbf{1} \\ & \quad x \geq \mathbf{0} \\ & \quad s \geq 1 \end{aligned} \tag{2}$$

The vector variable u is indexed by action sequences of Player 2. All optimal solutions to (2) agree on the value of this u -vector, which always takes the form of a 0/1-vector. The 1-entries in this 0/1-vector determine those inequalities of $P^{(0)}$ that may have *slack* in an optimal solution to $P^{(0)}$. Intuitively, they identify certain action sequences of Player 2 as *containing mistakes*. Let $\tilde{m}^{(1)}$ be this optimal u . We let $\tilde{M}^{(1)}$ be the set of sequences on which $\tilde{m}^{(1)}$ is 1. Let $m^{(1)}$ be defined by

$$m_{\sigma c}^{(1)} = \begin{cases} 1 & \text{if } \tilde{m}_\sigma^{(1)} = 0 \wedge \tilde{m}_{\sigma c}^{(1)} = 1 \\ 0 & \text{otherwise} \end{cases} \tag{3}$$

Intuitively, the 1-entries in this 0/1-vector identify certain actions of Player 2 as *being* mistakes, namely the final actions of the sequences on which $m^{(k)}$ is 1. We let $M^{(1)}$ be the set of actions c for which $m^{(1)}$ is 1 on the sequence ending in c . If $m^{(1)} \neq \mathbf{0}$, it defines the next round of linear programs.

Assuming we have already defined rounds $0, \dots, k-1$, the k 'th round looks as follows. We define the linear program $P^{(k)}$:

$$\begin{aligned}
& \underline{P^{(k)}} : \\
& \max_{x,q,t} && t \\
& \text{s.t.} && -A^\top x + F^\top q + m^{(k)}t \leq -\sum_{0 < i < k} m^{(i)}v^{(i)} \\
& && Ex = e \\
& && f^\top q = v^{(0)} \\
& && x \geq \mathbf{0} \\
& && t \geq 0
\end{aligned} \tag{4}$$

The variable t is scalar. We let $v^{(k)}$ be the value of t in an optimal solution of (4). Informally, $v^{(k)}$ is the *maximin slack* achievable in the inequalities indexed by $M^{(k)}$ (the min being over the inequalities). Note that $P^{(0)}$ is special and does not have the same format as $P^{(k)}$ for $k \geq 1$. Still, it is clear that the feasible solutions of $P^{(1)}$ are exactly the optimal solutions of $P^{(0)}$. In general, we will have that the feasible solutions of $P^{(k+1)}$ are the optimal solutions of $P^{(k)}$. Next, we consider the following linear program $Q^{(k)}$.

$$\begin{aligned}
& \underline{Q^{(k)}} : \\
& \max_{x,q,u,s} && \mathbf{1}^\top u \\
& \text{s.t.} && -A^\top x + F^\top q + u + \sum_{0 < i \leq k} m^{(i)}v^{(i)}s \leq \mathbf{0} \\
& && Ex - es = \mathbf{0} \\
& && f^\top q - v^{(0)}s = 0 \\
& && \mathbf{0} \leq u \leq \mathbf{1} \\
& && x \geq \mathbf{0} \\
& && s \geq 1
\end{aligned} \tag{5}$$

As was the case for $Q^{(0)}$, all optimal solutions to (5) agree on the value of the u -vector, which always takes the form of a 0/1-vector. The 1-entries in this 0/1-vector determine the inequalities of $P^{(k)}$ that may have slack in an optimal solution to $P^{(k)}$. Let $\tilde{m}^{(k+1)}$ be this optimal u . We let $\tilde{M}^{(k+1)}$ be the set of sequences on which $\tilde{m}^{(k+1)}$ is 1. Let $m^{(k+1)}$ be defined by

$$m_{\sigma c}^{(k+1)} = \begin{cases} 1 & \text{if } \tilde{m}_\sigma^{(k+1)} = 0 \wedge \tilde{m}_{\sigma c}^{(k+1)} = 1 \\ 0 & \text{otherwise} \end{cases} \tag{6}$$

We let $M^{(k+1)}$ be the set of actions c for which $m^{(k+1)}$ is 1 on the sequence ending in c . We denote this set as the *mistakes of order $k+1$* of Player 2. With this terminology, an informal interpretation of the meaning of variable $v^{(k)}$ is the *cost* for Player 2 of each mistake of order k that he includes in his chosen strategy. This interpretation is formalized in Lemma 14 in Appendix B. The sets $M^{(k)}$ are not necessarily disjoint, i.e., a specific action may be a mistake of several different orders. It can be seen that the set of orders of a given mistake is an interval $[i, i+1, \dots, j]$ but we shall not use this fact. This completes the description of the k 'th round. If $m^{(k+1)} \neq \mathbf{0}$, it defines the next round of linear programs.

In some round $k = K$, an optimal solution to (5) has $u = \mathbf{0}$ and the procedure is terminated. We argue in Appendix B (Lemma 12) that this happens after at most a number of rounds equal to the number of actions of Player 2 in the game tree. Let D be the x -parts of the set of optimal solutions to $P^{(K)}$. This is a set of realization plans for Player 1. Interchanging the role of Player 1 and Player 2 and negating the payoff matrix, we carry out the entire procedure again. Let D' be the resulting set of realization plans for Player 2. We can now state the main theorem of the present section.

Theorem 3 $D \times D'$ is the set of normal form proper equilibria of G in realization plans.

Since the number of rounds is at most the number of actions, standard arguments imply that the number of bits needed to describe the coefficients of the linear programs remains polynomially bounded throughout the iteration. Thus the theorem implies that we have a polynomial time algorithm for computing a normal form proper equilibrium in realization plans (and hence behavior plans) for a given two-player, zero-sum extensive form game of imperfect information but perfect recall. That is, we have established Theorem 1 of the introduction. However, one should notice that Theorem 3 allows us not only to compute *one* normal form proper equilibrium; it allows us to exhibit a linear program characterizing all of them. Due to lack of space the proof of Theorem 3 can be found in Appendix B.

To give a better intuition of how the algorithm works, we show the sequence of linear programs arising when finding the proper strategies for Player 1 in the game *Up or Down?* from Figure 3. The matrices defining the sequence form of the game are given below. All zero-entries are omitted for readability.

$$A = \begin{array}{ccccc} & \lambda & l & r & l' & r' \\ \left[\begin{array}{c} \lambda \\ U \\ D \end{array} \right] & & & 1 & & \end{array} \quad E = \begin{array}{ccc} & \lambda & U & D \\ \left[\begin{array}{c} 1 \\ -1 & 1 & 1 \end{array} \right] & & & \end{array} \quad F = \begin{array}{ccccc} & \lambda & l & r & l' & r' \\ \left[\begin{array}{c} 1 \\ -1 & 1 & 1 \\ -1 & & & 1 & 1 \end{array} \right] & & & & \end{array} \begin{array}{c} 0 \\ h \\ h' \end{array}$$

The first linear program to solve is $P^{(0)}$. Inserting the above matrices in (1), we get the linear program given below on the left. It describes the set of maximin strategies for Player 1, and the optimal value of the objective function is the value of the game. In this case, any optimal solution has $q_h = q_{h'} = q_0 = 0$ while any realization plan is an optimal choice of x . The optimal value of the objective function is $v^{(0)} = 0$. This value is used in the construction of $Q^{(0)}$, given below on the right.

$$\begin{array}{l} \underline{P^{(0)}} : \\ \max_{x,q} \\ \text{s.t.} \end{array} \quad \begin{array}{l} q_0 \\ q_0 \leq q_h + q_{h'} \\ q_h \leq 0 \\ q_h \leq x_U \\ q_{h'} \leq 0 \\ q_{h'} \leq 2x_D \\ x_\lambda = 1 \\ x_U + x_D = x_\lambda \\ x_\lambda, x_U, x_D \geq 0 \end{array} \quad \begin{array}{l} \underline{Q^{(0)}} : \\ \max_{x,q,u,s} u_\lambda + u_l + u_r + u_{l'} + u_{r'} \\ \text{s.t.} \end{array} \quad \begin{array}{l} q_0 + u_\lambda \leq q_h + q_{h'} \\ q_h + u_l \leq 0 \\ q_h + u_r \leq x_U \\ q_{h'} + u_{l'} \leq 0 \\ q_{h'} + u_{r'} \leq 2x_D \\ x_\lambda = s \\ x_U + x_D = x_\lambda \\ q_0 = 0 \\ 0 \leq u_\lambda, u_l, u_r, u_{l'}, u_{r'} \leq 1 \\ x_\lambda, x_U, x_D \geq 0 \\ s \geq 1 \end{array}$$

All optimal solutions to $Q^{(0)}$ have $u_{\{\lambda,l,l'\}} = \mathbf{0}$ and $u_{\{r,r'\}} = \mathbf{1}$, indicating that r and r' are the two possible mistakes of order 1 that Player 2 can make. This optimal u defines the vector $m^{(1)}$ of mistakes, which is used in the construction of $P^{(1)}$, given below on the left. This program maximizes the minimum slack in the inequalities corresponding to the mistakes found. In this case, the maximin slack achievable is $t = \frac{2}{3}$, achieved by setting $x_U = \frac{2}{3}$ and $x_D = \frac{1}{3}$. Intuitively, it is the cost for Player 2 of making either of his two possible mistakes. Inserting $v^{(1)} = \frac{2}{3}$ into (5), we obtain $Q^{(1)}$ given below on the right. All feasible solutions to $Q^{(1)}$ have $u = 0$, so there are no mistakes of order 2 for Player 1 to exploit. By Theorem 3, any optimal x to $P^{(1)}$ is therefore a proper strategy. The only one in this case is $x_U = \frac{2}{3}$, $x_D = \frac{1}{3}$.

$$\begin{array}{ll}
\underline{P^{(1)}} : & \\
\max_{x,q} & t \\
\text{s.t.} & q_0 \leq q_h + q_{h'} \\
& q_h \leq 0 \\
& q_h + t \leq x_U \\
& q_{h'} \leq 0 \\
& q_{h'} + t \leq 2x_D \\
& x_\lambda = 1 \\
& x_U + x_D = x_\lambda \\
& x_\lambda, x_U, x_D, t \geq 0 \\
& q_0 = 0 \\
& x_\lambda, x_U, x_D, t \geq 0
\end{array}
\qquad
\begin{array}{ll}
\underline{Q^{(1)}} : & \\
\max_{x,q,u,s} & u_\lambda + u_l + u_r + u_{l'} + u_{r'} \\
\text{s.t.} & q_0 + u_\lambda \leq q_h + q_{h'} \\
& q_h + u_l \leq 0 \\
& q_h + u_r \leq x_U - \frac{2}{3}s \\
& q_{h'} + u_{l'} \leq 0 \\
& q_{h'} + u_{r'} \leq 2x_D - \frac{2}{3}s \\
& x_\lambda = s \\
& x_U + x_D = x_\lambda \\
& q_0 = 0 \\
& 0 \leq u_\lambda, u_l, u_r, u_{l'}, u_{r'} \leq 1 \\
& x_\lambda, x_U, x_D \geq 0 \\
& s \geq 1
\end{array}$$

3 The procedure for perfect information games

Before we present our procedure for perfect information games, we want to address a subtle point. Our result for imperfect information games characterizes the normal form proper equilibria in realization plans which are equivalent to behavior *plans*. That is, no meaningful behavior is described in *irrelevant* information sets for the plans (See Appendix A for definitions). This is an inherent property of the notion of normal form proper equilibrium. However, it can be argued that it is sometimes relevant to prescribe meaningful behavior in irrelevant information sets. An example is the computation of prescriptive strategies for game playing software; one could imagine the software being used for advisory purposes with the user having the opportunity of ignoring the advice. Then, the user might still want meaningful advice after having ignored a particular piece of advice.

A slightly more restrictive concept than normal form properness that also provides meaningful behavior in irrelevant information sets was suggested by van Damme [16]. He calls an equilibrium in behavior strategies of an extensive form game *induced* by a normal form proper equilibrium if it is a limit of behavior strategies given by a sequence of ϵ -proper equilibria of the corresponding normal form. We shall adopt the slightly more convenient terminology *induced normal form proper equilibrium* for such an equilibrium. Note that while the induced normal form proper equilibrium is the *limit* of the behavior strategies *behaviorally equivalent* to the sequence of the ϵ -proper strategy profiles, the normal form proper equilibrium is the behavior strategy *behaviorally equivalent* to the *limit* of the sequence. It can be seen that the normal form proper equilibria are simply those equilibria that can be obtained by taking an induced normal form proper equilibrium and replacing the behavior in irrelevant information sets with arbitrary behavior. In particular, if we consider equilibria in behavior plans (as in the previous section) rather than behavior strategies the two notions coincide. But for perfect information games, we are able to characterize and compute the induced normal form proper equilibria in behavior *strategies*, rather than plans. It is an interesting open problem to do this for imperfect information games.

Let G be a perfect information zero-sum game played between Player 1, trying to maximize payoff and Player 2, trying to minimize payoff. The game is given by a game tree with payoffs in leaves and each internal node belonging to either Player 1, Player 2 or Chance. For each node i in the tree we associate three number $\underline{v}_i \leq v_i \leq \bar{v}_i$. The number v_i is the usual minimax value of the node and may be computed by standard backwards induction. The values \underline{v}_i and \bar{v}_i can be informally seen as pessimistic and optimistic estimates of the expected outcome of

the game from the point of Player 1, taking the possibility of mistakes being made by either player into account.

For a leaf with payoff p we let $\bar{v}_i = v_i = \underline{v}_i = p$. For an internal node i , we denote the set of immediate successors of i by $S(i)$ and define $\underline{v}_i, \bar{v}_i$ inductively as follows.

If i is a node belonging to Player 1, we let $V_i = (\cup_{j \in S(i)} \{v_j, \underline{v}_j\}) \setminus \{v_i\}$, i.e., the set of all values and all pessimistic estimates of all immediate successors of i , *except* the value of i itself. Then we let

$$\underline{v}_i = \begin{cases} \max(V_i) & \text{if } V_i \neq \emptyset, \\ v_i & \text{otherwise.} \end{cases} \quad (7)$$

Also, for a node i belonging to Player 1, we let $I_i = \{j \in S(i) | v_j = v_i \wedge \bar{v}_j > v_j\}$ and let

$$\bar{v}_i = v_i + \begin{cases} \frac{1}{\sum_{j \in I_i} (\bar{v}_j - v_j)^{-1}} & \text{if } I_i \neq \emptyset, \\ 0 & \text{otherwise.} \end{cases} \quad (8)$$

Similarly, if i is a node belonging to Player 2, we let $V_i = (\cup_{j \in S(i)} \{v_j, \bar{v}_j\}) \setminus \{v_i\}$, and let

$$\bar{v}_i = \begin{cases} \min(V_i) & \text{if } V_i \neq \emptyset, \\ v_i & \text{otherwise.} \end{cases} \quad (9)$$

Also, for node i belonging to Player 2, we let $I_i = \{j \in S(i) | v_j = v_i \wedge \underline{v}_j < v_j\}$ and let

$$\underline{v}_i = v_i - \begin{cases} \frac{1}{\sum_{j \in I_i} (v_j - \underline{v}_j)^{-1}} & \text{if } I_i \neq \emptyset, \\ 0 & \text{otherwise.} \end{cases} \quad (10)$$

If i is a node belonging to Chance and $j \in S(i)$ is chosen by Chance with probability α_j , we let $R_i = \{j \in S(i) | \underline{v}_j < v_j\}$ and $R'_i = \{j \in S(i) | \bar{v}_j > v_j\}$ and let

$$\underline{v}_i = v_i - \min_{j \in R_i} \alpha_j (v_j - \underline{v}_j) \quad (11)$$

$$\bar{v}_i = v_i + \min_{j \in R'_i} \alpha_j (\bar{v}_j - v_j) \quad (12)$$

Our characterization of induced proper equilibria is then the following:

Theorem 4 *A behavior strategy profile ρ for G is an induced normal form proper equilibrium of G if and only if the following three conditions all hold:*

1. *For all nodes i and immediate successors j , ρ assigns non-zero behavior probability to j only if $v_i = v_j$.*
2. *For all nodes i belonging to Player 1 for which $I_i \neq \emptyset$, ρ assigns behavior probability exactly $\frac{(\bar{v}_j - v_j)^{-1}}{\sum_{j \in I_i} (\bar{v}_j - v_j)^{-1}}$ to each $j \in I_i$.*
3. *For all nodes i belonging to Player 2 for which $I_i \neq \emptyset$, ρ assigns behavior probability exactly $\frac{(v_j - \underline{v}_j)^{-1}}{\sum_{j \in I_i} (v_j - \underline{v}_j)^{-1}}$ to each $j \in I_i$.*

Note that the first condition is exactly the condition of being a subgame perfect equilibrium, while the other conditions further restrict the behavior probabilities of choices made in nodes i for which $I_i \neq \emptyset$, to uniquely defined values. The theorem immediately implies Theorem 2 of the introduction. Due to lack of space, the proof can be found in Appendix C.

References

- [1] Darse Billings, Neil Burch, Aaron Davidson, Robert Holte, Jonathan Schaeffer, Terence Schauenberg, and Duane Szafron. Approximating game-theoretic optimal strategies for full-scale poker. In *International Joint Conference on Artificial Intelligence*, 2003.
- [2] Xi Chen and Xiaotie Deng. Settling the complexity of two-player Nash equilibrium. In *Proceedings of 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*, pages 261–272, 2006.
- [3] Melvin Dresher. *The Mathematics of Games of Strategy: Theory and Applications*. Prentice-Hall, 1961.
- [4] Andrew Gilpin and Tuomas Sandholm. Finding equilibria in large sequential games of incomplete information. Technical Report CMU-CS-05-158, Carnegie Mellon University, 2005.
- [5] Elon Kohlberg and Jean-Francois Mertens. On the strategic stability of equilibria. *Econometrica*, 54:1003–1037, 1986.
- [6] Daphne Koller, Nimrod Megiddo, and Bernhard von Stengel. Fast algorithms for finding randomized strategies in game trees. In *Proceedings of the 26th Annual ACM Symposium on the Theory of Computing*, pages 750–759, 1994.
- [7] Daphne Koller, Nimrod Megiddo, and Bernhard von Stengel. Efficient computation of equilibria for extensive form games. *Games and Economic Behavior*, 14:247–259, 1996.
- [8] Daphne Koller and Avi Pfeffer. Representations and solutions for game-theoretic problems. *Artificial Intelligence*, 94(1–2):167–215, 1997.
- [9] David M. Kreps and Robert Wilson. Sequential equilibria. *Econometrica*, 50(4):863–894, July 1982.
- [10] H.W. Kuhn. Extensive games and the problem of information. *Annals of Mathematical Studies*, 28:193–216, 1953.
- [11] Richard D. McKelvey, Andrew M. McLennan, and Theodore L. Turocy. Gambit: Software tools for game theory, version 0.97.0.7. <http://econweb.tamu.edu/gambit>, 2004.
- [12] Peter Bro Miltersen and Troels Bjerre Sørensen. Computing proper equilibria of zero-sum games. In *Proceedings of the fifth International Conference on Computers and Games*, May 2006.
- [13] Peter Bro Miltersen and Troels Bjerre Sørensen. Computing sequential equilibria for two-player games. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 107–116, 2006.
- [14] Peter Bro Miltersen and Troels Bjerre Sørensen. Computing a quasi-perfect equilibrium of a two-player game. Invited to special issue of *Economic Theory* on equilibrium computation, to appear, 2007.
- [15] R. B. Myerson. Refinements of the Nash equilibrium concept. *International Journal of Game Theory*, 15:133–154, 1978.

- [16] Eric van Damme. A relation between perfect equilibria in extensive form games and proper equilibria in normal form games. *International Journal of Game Theory*, 13:1–13, 1984.
- [17] Eric van Damme. *Stability and Perfection of Nash Equilibria*. Springer-Verlag, 2nd edition, 1991.
- [18] Bernhard von Stengel. Efficient computation of behavior strategies. *Games and Economic Behavior*, 14:220–246, 1996.
- [19] Yoshitsugu Yamamoto. A path-following procedure to find a proper equilibrium of finite games. *International Journal of Game Theory*, 22:249–259, 1993.

Appendix A

In this appendix, we describe the main concepts of extensive-form games that are used in this paper, and in particular in the proofs of Appendix B and Appendix C. For more details on extensive form games in general, see any textbook on game theory and for more details on sequence form representation, see von Stengel [18] and Koller, Megiddo and von Stengel [7].

A two-player, extensive-form game zero-sum game G with imperfect information but perfect recall is given by a finite tree with payoffs at the leaves, information sets partitioning nodes of the tree and nodes of chance being allowed. Actions of a player are denoted by labels on edges of the tree. Perfect recall means that all nodes in an information set belonging to a player share the sequence of actions and information sets of that player that are visited on the path from the root to the nodes. For convenience, we assume that actions taken in different information sets have different names. In particular, from the name of an action, we can deduce its information set. Also, by perfect recall, from the name c of an action, we can deduce the sequence of actions σ taken before c by the same player. We shall use the following notation in these appendices.

Definition 5 *Given an information set h , we let C_h be the set of actions that may be taken in this information set. Given a sequence of actions σ of a player (starting with his first action in the play), we let S_σ be the set of information sets in which the player may be required to take an action immediately following σ .*

Note that S_σ is empty if the player is guaranteed not to be required to take any further actions following σ .

A *pure strategy* s for a player is a *set* of designated actions, containing exactly one action of each information set belonging to him. A *mixed strategy* for a player is a probability distribution over pure strategies. Since there are exponentially many pure strategies as a function of the number of information sets, it is usually not feasible to explicitly represent mixed strategies. A more compact notion is the notion of a *behavior strategy*. A behavior strategy is simply an assignment of probabilities to actions, and therefore an object of size comparable to the description of the game itself. Kuhn [10] showed that for games of perfect recall, for any mixed strategy p , there is a behavior strategy π that is *behaviorally equivalent*. That is, a player playing by p will generate exactly the same distribution of plays against any opponent as a player playing by π . Thus, we can compactly represent mixed strategies by behavior strategies. An even more compact notion is the notion of a *behavior plan*. This is an assignment of probabilities to actions of those information sets *that may be reached with non-zero probability if the plan is played* against some strategy of the opponent. These information sets are called *relevant* for the plan. No meaningful behavior is assigned to information sets that are guaranteed not to be reached if the plan is played, no matter what the opponent does. These information sets are called *irrelevant* for the plan. Behavior strategies and behavior plans are clearly behaviorally equivalent.

An important insight of von Stengel [18] and Koller, Megiddo and von Stengel is that behavior plans for games of perfect recall are for computational purposes often better represented in *sequence form*, which we describe next. Given a behavior strategy for one of the players, the *realization weight* of a sequence σ of actions made by that player is the *product* of behavior probabilities of the actions in the sequence. A *realization plan* for a player is a vector of realization weights indexed by his possible sequences of actions. Note that a behavior probability is the ratio between two realization weights. If this ratio is 0/0, the realization weights do not define a behavior probability. Still, the map between behavior *plans* and realization plans is a bijection, so given a realization plan, we may talk about the corresponding behavior plan and vice versa. Note that a pure strategy corresponds to a realization plan that is a 0/1-vector. For a game G

of perfect recall, the statement that a vector is a valid realization plan can be expressed by a non-negativity constraint and a system of linear equations indexed by information sets. We let

$$Ex = e, x \geq 0 \quad (13)$$

denote the constraints expressing that x is a valid realization plan for Player 1 and we let

$$Fy = f, y \geq 0 \quad (14)$$

be the constraints expressing that y is a valid realization plan for Player 2. Here, E and F are matrices containing entries from $\{-1, 0, 1\}$. The rows of E (F) are indexed by information sets of Player 1 (Player 2) and the vector e (f) is the vector with a 1 in the entry indexed by the information set in the root of the game and is 0 elsewhere. Here, for convenience, we assume that the root of the game is a singleton information set belonging to both players.

The key to the computational usefulness of sequence form is the following fact: If Player 1 plays by realization plan x and Player 2 plays by realization plan y , then the expected payoff for Player 1 is given by a *bilinear form* $x^\top Ay$, where A is easily computed from the description of the game. Koller, Megiddo and von Stengel used this to show that the set of maximin realization plans for Player 1 is given by the x -parts of optimal solutions to the following linear program.

$$\begin{aligned} \max_{x, q} \quad & f^\top q \\ \text{s.t.} \quad & -A^\top x + F^\top q \leq 0 \\ & Ex = e \\ & x \geq \mathbf{0} \end{aligned} \quad (15)$$

The value of the objective function in an optimal solution is the value of the game. Note that $f^\top q$ is simply q_0 where 0 is the information set at the root. The inequalities in the system $-A^\top x + F^\top q \leq 0$ are indexed by action sequences of Player 2. For understanding how the program works, it is very useful to know the following interpretation of the variable q_h : Player 2 can guarantee that the contribution to the expected payoff from plays where he finds himself in information set h is at most q_h . This gives the inequalities of $-A^\top x + F^\top q \leq 0$ a very natural interpretation. Consider an action c . There is a unique information set h in which c is taken and a unique sequence σ taken before c . Then, the inequality associated with the sequence σc may be rearranged to the following.

$$q_h \leq (A^\top x)_{\sigma c} + \sum_{h' \in S_{\sigma c}} q_{h'}. \quad (16)$$

Here, the contribution from the first term is from those cases where the game ends before Player 2 is required to take another action after c . The contribution from the remaining terms are from those cases where he finds himself required to take further action. Given the above interpretation of q_h , the inequality is very natural. It states that Player 2 is no worse off than he would be if he had to commit himself to the action c . From the fact that all inequalities of (15) have the form of (16) we get the following fact, central for the reasoning in Appendix B.

Fact 6 *Given a fixed choice of x in (15), an optimum feasible choice of q can be computed inductively as follows:*

$$q_h = \min_{c \in C_h} \left((A^\top x)_{\sigma c} + \sum_{h' \in S_{\sigma c}} q_{h'} \right). \quad (17)$$

Appendix B

In this appendix, we prove Theorem 3. The main idea is similar to the proof of van Damme's characterization of the proper equilibria of a matrix game [17, Theorem 3.5.5]. We prove that a proper strategy will be an optimal solution to all the linear programs $P^{(k)}$ (Lemma 16 below) and we prove that any two optimal solutions to the final program $P^{(K)}$ are payoff-equivalent (follows from Lemma 14 below). Theorem 3 then follows easily from these two facts. But first, we have to prove a series of lemmas partially characterizing the structure of feasible and optimal solutions to our linear programs and the structure of the sets $\tilde{M}^{(k)}$ and $M^{(k)}$. The proof of the first lemma (stated for convenience) is obvious.

Lemma 7 *The optimal solutions to $P^{(k)}$ are the feasible solutions to the following set of equations and inequalities, $R^{(k)}$:*

$$\begin{aligned} -A^\top x + F^\top q &\leq -\sum_{0 < i \leq k} m^{(i)} v^{(i)} \\ Ex &= e \\ f^\top q &= v^{(0)} \\ x &\geq \mathbf{0} \end{aligned} \tag{18}$$

The following lemma explains the purpose of $Q^{(k)}$: It identifies the sequences whose corresponding inequality may be satisfied with *slack* (i.e., as a strict inequality) in an optimal solution to $P^{(k)}$ (i.e., a feasible solution to $R^{(k)}$).

Lemma 8 *The vector $\tilde{m}^{(k)}$ is indeed a $\{0, 1\}$ vector, so $\tilde{M}^{(k)}$ is well-defined. A sequence σ is in $\tilde{M}^{(k)}$ iff there is a feasible solution to $R^{(k-1)}$ with slack in the inequality indexed by σ in the system $-A^\top x + F^\top q \leq -\sum_{0 < i < k} m^{(i)} v^{(i)}$. Furthermore, there is a single solution with slack in all the inequalities indexed by $\tilde{M}^{(k)}$.*

Proof The linear program $Q^{(k)}$ used to identify $\tilde{m}^{(k)}$ is a straightforward application of a general technique for identifying slack in systems of linear inequalities. Given a set of linear inequalities with a non-empty solution set

$$\begin{aligned} A^\top x &\leq b \\ x &\geq \mathbf{0}, \end{aligned} \tag{19}$$

we need to prove that the following linear program identifies the inequalities of (19) where slack is achievable:

$$\begin{aligned} \max_{x, s, u} \quad & \mathbf{1}^\top u \\ \text{s.t.} \quad & A^\top x - bs + u \leq \mathbf{0} \\ & \mathbf{0} \leq u \leq \mathbf{1} \\ & s \geq 1 \\ & x \geq \mathbf{0} \end{aligned} \tag{20}$$

Since the set defined by (19) is convex, there is a solution that has slack in all the inequalities where slack is achievable: We may simply take a convex combinations of the solutions achieving slack in each particular inequality. Given such a solution x to (19), we construct a solution (x^*, s^*, u^*) to (20) with objective value equal to the number of inequalities with slack: Let δ be the smallest strictly positive slack in (19) for the given x . Then let $s^* = \max\{1, 1/\delta\}$, $x^* = xs^*$ and u be the $\{0, 1\}$ -vector indicating the inequalities with slack in (19). This is a *scaling* of (19), such that even the smallest positive slack is now at least 1. Since u is only 1 where there

was already a slack, the solution (x^*, s^*, u^*) is feasible, and clearly has objective value equal to the number of inequalities with slack in (19). It is also easy to see that if a given 0/1 vector u^* is achievable in an optimal solution, the vector x^* of that solution is a scaling of a solution to (19) with slack in the corresponding inequalities. ♠

Lemma 9 *The (x, q) -parts of the feasible solutions to $P^{(k+1)}$ are exactly the (x, q) -parts of the optimal solutions to $P^{(k)}$.*

Proof The statement follows directly from Lemma 7 and Lemma 8. ♠

Lemma 10 *For any action sequence $\sigma \in \tilde{M}^{(k)}$ with the extension σc being also a legal action sequence, we have $\sigma c \in \tilde{M}^{(k)}$.*

Proof Recall that the inequalities $-A^\top x + F^\top q \leq -\sum_{0 < i < k} m^{(i)} v^{(i)}$ of $R^{(k-1)}$ are indexed by Player 2's action sequences. That $\sigma \in \tilde{M}^{(k)}$ means that $\tilde{m}_\sigma = 1$ and, by Lemma 8, that it is possible to achieve slack in the inequality indexed by σ in $-A^\top x + F^\top q \leq -\sum_{0 < i < k} m^{(i)} v^{(i)}$ by a feasible solution to $R^{(k-1)}$. This inequality is the inequality

$$q_h \leq (A^\top x)_\sigma + \sum_{h' \in S_\sigma} q_{h'} - \sum_{j < k} m_\sigma^{(j)} v^{(j)} \quad (21)$$

where h is the information set where the last action of σ is chosen. That is, we have a feasible solution (\bar{x}, \bar{q}) to $R^{(k-1)}$ with

$$\bar{q}_h = (A^\top x)_\sigma + \sum_{h' \in S_\sigma} \bar{q}_{h'} - t - \sum_{j < k} m_\sigma^{(j)} v^{(j)} \quad (22)$$

for some $t > 0$. We need to show that we can find a feasible solution with slack in the inequality

$$q_{\bar{h}} \leq (A^\top x)_{\sigma c} + \sum_{h' \in S_{\sigma c}} q_{h'} - \sum_{j < k} m_{\sigma c}^{(j)} v^{(j)} \quad (23)$$

where \bar{h} is the information set where the action c may be chosen. Note that $\bar{h} \in S_\sigma$. We may obtain such a feasible solution simply subtracting t from the value of $q_{\bar{h}}$ in the current feasible solution while leaving the value of all other variables unchanged. The new solution is still feasible (as \bar{h} is not in S_τ for any $\tau \neq \sigma$, by perfect recall) and the slack has been transferred from inequality (21) to inequality (23). ♠

Lemma 11 *For fixed k and on any path from the root to a leaf in the game tree, at most one action of Player 2 is a member of $M^{(k)}$.*

Proof An action c taken after σ is in $M^{(k)}$ if and only if σc is in $\tilde{M}^{(k)}$ but σ is not. By lemma 10, if a given sequence is in $\tilde{M}^{(k)}$, all extensions of that sequence is also in $\tilde{M}^{(k)}$. Thus, if c is in $M^{(k)}$, no action taken after c is in $M^{(k)}$. ♠

Lemma 12 *For any k , $\tilde{M}^{(k+1)}$ is a proper subset of $\tilde{M}^{(k)}$.*

Proof By lemma 8, $\tilde{M}^{(k+1)}$ is the set of indices of the inequalities of $R^{(k)}$ that may be satisfied with slack by some feasible solution, while $\tilde{M}^{(k)}$ is the corresponding set of indices of $R^{(k-1)}$. Since $R^{(k)}$ is obtained from $R^{(k-1)}$ by adding a negative-valued term to the right hand sides of some of the inequalities, we have that $\tilde{M}^{(k+1)} \subseteq \tilde{M}^{(k)}$.

To see that the inclusion is strict, assume to the contrary that $\tilde{M}^{(k+1)} = \tilde{M}^{(k)}$. Then all inequalities that can be satisfied with slack in $R^{(k-1)}$ can also be satisfied with slack in $R^{(k)}$. But observe that by construction of the program $P^{(k)}$ and Lemma 8, the value $v^{(k)}$ is the largest possible minimum slack by a feasible solution to the inequalities of $R^{(k-1)}$, in those inequalities indexed by $M^{(k)}$. If $\tilde{M}^{(k)} = \tilde{M}^{(k+1)}$, we also have $M^{(k)} = M^{(k+1)}$ and if all inequalities indexed by $M^{(k)}$ are also indexed by $M^{(k+1)}$, a larger minimum slack than $v^{(k)}$ were in fact possible in solutions to $R^{(k-1)}$; a contradiction. ♠

Note that Lemma 12 implies that the procedure terminates after at most r rounds, where r is the total number of actions of Player 2.

The following lemma is a central observation. It captures the fact that the variables q_h have exactly the same semantics in $P^{(k)}$ as in the original Koller-Megiddo-von Stengel program (see Appendix A). It also implies that when given the x -part of an optimal solution to $P^{(k)}$ only, the q and t -part can be computed in linear time in the size of the game tree by the obvious “backwards induction”.

Lemma 13 *There are optimal solutions to $P^{(k)}$ with $q = \bar{q}$ of the following form:*

$$\bar{q}_h = \min_{c \in C_h} \left((A^\top x)_{\sigma c} + \sum_{h' \in S_{\sigma c}} \bar{q}_{h'} \right) \quad (24)$$

Proof Note that a priori, it is neither obvious that the stated \bar{q} is feasible, nor that it is optimal. We shall establish by induction in k that it is feasible as well as optimal.

Fix optimal x . For $k = 0$, the stated \bar{q} is feasible and optimal by Fact 6 in the discussion of the Koller-Megiddo-von Stengel program (Appendix A). Now assume $k \geq 1$ and assume that the stated \bar{q} is feasible and optimal for $P^{(k-1)}$. By Lemma 9, it is feasible for $P^{(k)}$, so we only need to show that it is also optimal.

For this, let \bar{t} be the minimum slack in the inequalities indexed by $M^{(k)}$ in the solution defined by \bar{q} , so we have that the optimal value of the objective function for our choice of \bar{q} is \bar{t} . Let σc be the sequence indexing the inequality of $R^{(k-1)}$ obtaining this minimum slack \bar{t} in the solution (x, \bar{q}, \bar{t}) . Let h be the information set in which the choice c is taken. That is, we have

$$\bar{q}_h = (A^\top x)_{\sigma c} + \sum_{h' \in S_{\sigma c}} \bar{q}_{h'} - \bar{t} - \sum_{j < k} m_{\sigma c}^{(j)} v^{(j)}. \quad (25)$$

With $\sigma = c_1 c_2 \dots c_l$, we have by Lemma 11 that $c_i \notin M^{(k)}$ for $i = 1, \dots, l$. This means that for those information sets h_i of Player 2 where the actions c_i are chosen, there is no slack in the inequalities

$$q_{h_i} \leq (A^\top x)_{c_1 c_2 \dots c_i} + \sum_{h' \in S_{c_1 c_2 \dots c_i}} q_{h'} - \sum_{j < k} m_{c_1 c_2 \dots c_i}^{(j)} v^{(j)} \quad (26)$$

of the system $R^{(k-1)}$. That is, we in fact have

$$\bar{q}_{h_i} = (A^\top x)_{c_1 c_2 \dots c_i} + \sum_{h' \in S_{c_1 c_2 \dots c_i}} \bar{q}_{h'} - \sum_{j < k} m_{c_1 c_2 \dots c_i}^{(j)} v^{(j)} \quad (27)$$

for all i . Now assume some other feasible (q^*, t^*) has $t^* > \bar{t}$. We can assume without loss of generality that t^* is the minimum slack in the inequalities indexed by $M^{(k)}$. Since $t^* > \bar{t}$ we must have that

$$q_h^* < (A^\top x)_{\sigma c} + \sum_{h' \in S_{\sigma c}} q_{h'}^* - \bar{t} - \sum_{j < k} m_{\sigma c}^{(j)} v^{(j)} \quad (28)$$

Note that \bar{q} has, by construction, the largest possible feasible value of q at every entry (it is defined as the minimum of a set of values, and there are constraints in $R^{(k)}$ stating that it must be less than or equal to all of them). Thus, we in fact have that

$$q_h^* < (A^\top x)_{\sigma c} + \sum_{h' \in S_{\sigma c}} \bar{q}_{h'} - \bar{t} - \sum_{j < k} m_{\sigma c}^{(j)} v^{(j)} = \bar{q}_h. \quad (29)$$

This means that $q_h^* < \bar{q}_h$. Now, by the fact that (27) holds for all i and the fact that q^* must satisfy the inequalities (26) of $R^{(k-1)}$ for all i , we have by induction that $q_{h_i}^* < \bar{q}_{h_i}$ for all i . In particular, we have $q_{h_0}^* < \bar{q}_{h_0}$. However, h_0 is the root information set of Player 2, so $q_{h_0}^* = f^\top q^* = v^{(0)}$ and $\bar{q}_{h_0} = f^\top \bar{q} = v^{(0)}$, a contradiction. \spadesuit

Lemma 13 allows for some very convenient terminology. When we below talk about some x being an optimal solution to $P^{(k)}$, we assume without loss of generality that q is being set as specified by Lemma 13 and t set in the obvious optimal way as the minimum slack in the inequalities indexed by $M^{(k)}$.

Lemma 14 *Let x be an optimal solution to $P^{(k)}$. Let s be a pure strategy of Player 2 and let y be the corresponding realization plan. For $j \leq k$, let $d^{(j)}$ be the number of actions prescribed by s that are also in $M^{(j)}$. Then,*

$$x^\top Ay \geq v^{(0)} + \sum_{j=1}^k d^{(j)} v^{(j)}. \quad (30)$$

If s has the property that it does not prescribe any actions from $M^{(k+1)}$, we have in fact that

$$x^\top Ay = v^{(0)} + \sum_{j=1}^k d^{(j)} v^{(j)}. \quad (31)$$

Proof By Lemma 7, x is a feasible solution to $R^{(k)}$. In particular, x satisfies the inequalities $-A^\top x + F^\top q \leq -\sum_{i=1}^k v^{(i)} m^{(i)}$. Taking inner product of y and $-A^\top x + F^\top q$, we get

$$y^\top (-A^\top x + F^\top q) \leq y^\top \left(-\sum_{i=1}^k v^{(i)} m^{(i)} \right) \quad (32)$$

Removing the parenthesis and transposing:

$$-x^\top Ay + q^\top Fy \leq -\sum_{i=1}^k v^{(i)} m^{(i)\top} y \quad (33)$$

Rearranging:

$$x^\top Ay \geq q^\top Fy + \sum_{i=1}^k v^{(i)} m^{(i)\top} y \quad (34)$$

$$= q^\top f + \sum_{i=1}^k v^{(i)} d^{(i)} \quad (35)$$

$$= v^{(0)} + \sum_{i=1}^k v^{(i)} d^{(i)} \quad (36)$$

Also, by Lemma 8, the inequalities of $-A^\top x + F^\top q \leq -\sum_{i=1}^k v^{(i)} m^{(i)}$ are in fact equalities, except on those inequalities indexed by an element of $\tilde{M}^{(k+1)}$. Now assuming further than s does not use an action of $M^{(k+1)}$, then $y_\sigma = 0$ for $\sigma \in \tilde{M}^{(k+1)}$. This turns all inequalities in the above derivation into equalities. \spadesuit

Lemma 15 *Let $k \geq 1$. If \bar{x} is a feasible but not optimal solution to $P^{(k)}$, then there is some pure strategy s of Player 2 prescribing some action of $M^{(k)}$, so that $\bar{x}^\top Ay < \sum_{i=0}^k v^{(i)}$, where y is the realization plan corresponding to s .*

Proof Since \bar{x} is a feasible solution to $P^{(k)}$, by Lemma 9, it is an optimal solution to $P^{(k-1)}$. Furthermore, we may assume that it is an optimal solution to $P^{(k-1)}$ using the generic \bar{q} of Lemma 13. Since we assume that \bar{x} is not an optimal solution to $P^{(k)}$, some inequality of $R^{(k)}$ is violated. That is, we can find a sequence $\bar{\sigma}\bar{c}$ where the choice \bar{c} is taken in an information set \bar{h} so that

$$(-A^\top \bar{x} + F^\top \bar{q})_{\bar{\sigma}\bar{c}} > (-\sum_{i=1}^k v^{(i)} m^{(i)})_{\bar{\sigma}\bar{c}}. \quad (37)$$

Let $\Delta = (-A^\top \bar{x} + F^\top \bar{q})_{\bar{\sigma}\bar{c}} - (-\sum_{i=1}^k v^{(i)} m^{(i)})_{\bar{\sigma}\bar{c}}$. Let s be the pure strategy that contains the sequence $\bar{\sigma}\bar{c}$, i.e., chooses all actions in $\bar{\sigma}\bar{c}$ in the information sets where this sequence is played, and in every other information set h chooses the action

$$\operatorname{argmin}_{c \in C_h} \left((A^\top \bar{x})_{\sigma c} + \sum_{h' \in S_{\sigma c}} \bar{q}_{h'} \right) \quad (38)$$

Note that s is a strategy that contains exactly one action from each of the sets $M^{(j)}$, $j \leq k$. Let y be the realization plan corresponding to this strategy. For the realization plan x , we have that $(-A^\top \bar{x} + F^\top \bar{q})_{\bar{\sigma}\bar{c}} = \Delta + (-\sum_{i=1}^k v^{(i)} m^{(i)})_{\bar{\sigma}\bar{c}}$ and $(-A^\top \bar{x} + F^\top \bar{q})_\sigma = (-\sum_{i=1}^k v^{(i)} m^{(i)})_\sigma$ on all other sequences σ where y_σ is 1. A similar calculation as in the proof of Lemma 14 now yields $\bar{x}^\top Ay = v^{(0)} + \sum_{i=1}^k v^{(i)} - \Delta$. \spadesuit

Lemma 16 *Let (\bar{x}, \bar{y}) be a normal form proper equilibrium of G in realization plans. Then, for all k , \bar{x} is an optimal solution to $P^{(k)}$.*

Proof Assume not. By the definition of proper equilibrium, there is a family p_ϵ of ϵ -proper fully mixed strategy profiles converging as $\epsilon \rightarrow 0$ to a mixed strategy profile p to which (\bar{x}, \bar{y}) is behaviorally equivalent. By van Damme [17, Lemma 2.3.2], for all sufficiently small ϵ , \bar{x} is a best response to the realization plan y_ϵ of Player 2 given by p_ϵ . We shall arrive at a contradiction by exhibiting a better response to y_ϵ than \bar{x} .

If \bar{x} is not an optimal solution to $P^{(k)}$, we can let j be the smallest j so that \bar{x} is not an optimal solution to $P^{(j)}$ but \bar{x} is an optimal solution to $P^{(i)}$ for all $i < j$. A proper equilibrium is also a Nash equilibrium, so $j \geq 1$. By Lemma 15, we can find a pure realization plan y of Player 2 so that $\bar{x}^\top Ay < \sum_{i=0}^j v^{(i)}$. Now, let x^* be an optimal solution to $P^{(j)}$. We shall argue that x^* is a better response to y_ϵ than \bar{x} is, and we shall be done.

To see that x^* is a better response to y_ϵ than \bar{x} is, we divide the pure strategies of Player 2 into three categories.

- (a) Those pure strategies s not using an action in $M^{(j)}$. By Lemma 14, \bar{x} and x^* has the same payoff against all of these.

- (b) Those pure strategies s that uses an action in $M^{(j)}$ and with the property that $\bar{x}^\top Ay < \sum_{i=0}^j v^{(i)}$, where y is the realization plan corresponding to s . We already identified one such s ; there may be others.
- (c) Those pure strategies s using an action in $M^{(j)}$ and with the property that $\bar{x}^\top Ay \geq \sum_{i=0}^j v^{(i)}$, where y is the realization plan corresponding to s .

The expected payoff of x^* (resp., \bar{x}) against y_ϵ is the weighted sum of the payoff of x^* (resp., \bar{x}) against each s , weighted by $p_\epsilon(s)$. Strategies of type (a) makes the same contribution to the payoff of x^* as to the payoff of \bar{x} . Strategies of type (b) puts x^* ahead of \bar{x} . Strategies of type (c) may go either way, but by the definition of an ϵ -proper strategy profile, we have for any s of type (c) and any s' of type (b) and sufficiently small ϵ that $p_\epsilon(s) \leq \epsilon p_\epsilon(s')$. Since there is at least one strategy of type (b), for sufficiently small ϵ , the contributions from all strategies of type (c) to the two weighted sums are negligible, compared to the contribution from strategies of type (b). Therefore, x^* is a better response to y_ϵ than \bar{x} is, and we have a contradiction. ♠

We are now ready to show the final lemma of this appendix that immediately implies Theorem 3.

Lemma 17 *Let K be the last round in our sequence of pairs of linear programs, i.e., the optimal solution to $Q^{(K)}$ has $u = \mathbf{0}$. Then the proper strategies of Player 1 in realization plan representation are exactly the x -parts of the optimal solutions to $P^{(K)}$.*

Proof Lemma 16 implies that realization plans corresponding to all proper strategies of Player 1 can be found among the optimal solutions to $P^{(K)}$ where K is the index of the last round. Also, Myerson's existence proof implies that at least one such proper equilibrium exists. Lemma 14 implies that for any two optimal solutions \bar{x} and x^* to $P^{(K)}$ and *any* realization plan y of Player 2, we have $\bar{x}^\top Ay = x^{*\top} Ay$, since no actions are mistakes of order $K + 1$, round K being the last round. We thus merely have to show the intuitively obvious fact that if \bar{x} is the realization plan corresponding to a proper strategy and x^* is payoff-equivalent to \bar{x} in this strong way, then x^* is also the realization plan corresponding to a proper strategy. For this, we may argue as van Damme [17, Page 61]: If p_ϵ is a sequence of ϵ -proper strategy profiles converging to a mixed strategy behaviorally equivalent to \bar{x} , we define

$$p'_\epsilon = \epsilon p_\epsilon + (1 - \epsilon)q \tag{39}$$

where q is a mixed strategy behaviorally equivalent to x^* . Then, it is easy to see that for sufficiently small ϵ , p'_ϵ is a sequence of ϵ -proper strategy profiles converging to q and hence x^* is a proper strategy in realization plan representation. ♠

Appendix C

In this appendix we prove Theorem 4. As matrix games play a more explicit role in this proof than in the rest of the paper, we introduce some explicit notation for them. A matrix game is a map $G : S_1^G \times S_2^G \rightarrow \mathfrak{R}$, where S_1^G is the set of pure strategies for Player 1 trying to maximize payoff, S_2^G is the set of pure strategies for Player 2 trying to minimize payoff and $G(i, j)$ is the payoff if Player 1 plays i and Player 2 plays j . The map G is extended to mixed strategies in the natural way, i.e., $G(p_1, p_2)$ is the expected payoff if Player 1 plays mixed strategy p_1 and Player 2 plays mixed strategy p_2 . For matrix games, the set of Nash equilibria is a Cartesian product $M_1^G \times M_2^G$ where M_1 is the set of maximin mixed strategies for Player 1 while M_2 is the set of minimax mixed strategies for Player 2. The value v_G of the game is the expected payoff for Player 1 if any equilibrium strategy is played, i.e., $v_G = G(p_1, p_2)$ for all $p_1 \in M_1^G$ and all $p_2 \in M_2^G$.

Given a matrix game G , the *essential* strategies $E_i^G \subseteq S_i^G$ are those strategies played with positive probability in some strategy in M_i^G while the *superfluous* strategies $F_i^G = S_i^G - E_i^G$ are those strategies played with probability 0 in all strategies in M_i^G . Given a matrix game G for which $F_2^G \neq \emptyset$, we shall consider a derived game \bar{G} defined as follows: Player 1 chooses a (possibly mixed) maximin strategy for Player 1 in G and Player 2 simultaneously chooses a pure strategy from the superfluous strategies F_2^G of G . The payoff is the expected payoff if the two strategies are played against each other in G . Since the set of strategies for Player 1 is convex and the payoff function is a linear function in his strategy, Player 1 has a pure maximin strategy in \bar{G} which is by definition a mixed strategy in G . The set of maximin strategies $M_1^{\bar{G}}$ for Player 1 in \bar{G} can therefore be regarded as a subset of M_1^G . Completely analogously, given a matrix game G for which $F_1^G \neq \emptyset$, we shall consider a derived game \underline{G} defined as follows: Player 2 chooses a (possibly mixed) minimax strategy for Player 2 in G and Player 1 simultaneously chooses a pure strategy from the superfluous strategies F_1^G of G . The payoff is the expected payoff if the two strategies are played against each other in G . The set of minimax strategies $M_2^{\underline{G}}$ for Player 2 in \underline{G} can be regarded as a subset of M_2^G . We now state two theorems, both proved by van Damme:

Theorem 18 [17, Theorem 3.5.5] *The set of proper equilibria of a matrix game G is a Cartesian product $D_1^G \times D_2^G$. Furthermore, if $F_1^G \neq \emptyset$, we have that $D_1^G \subseteq M_1^{\bar{G}}$ and if $F_1^G = \emptyset$, we have that $D_1^G = M_1^G$. Analogously, if $F_2^G \neq \emptyset$, we have that $D_2^G \subseteq M_2^{\underline{G}}$ and if $F_2^G = \emptyset$, we have that $D_2^G = M_2^G$.*

Theorem 19 [16, Theorem 1] *Let G be a finite two-player extensive form game with perfect recall. Let π be an induced normal form proper equilibrium of G . Let G' be a subgame of G . Let π' be the behavior strategy profile π restricted to G' . Then π' is an induced normal form proper equilibrium of G' .*

Actually, the statement of van Damme's Theorem 1 is somewhat weaker, but his proof immediately yields the above. Note that the statement is far from obvious. In particular, it is false for Nash equilibria and it is also false for (non-induced) normal form proper equilibria whose behavior in irrelevant subgames can be arbitrary.

Using the above theorems, we now prove Theorem 4 by structural induction in the tree representing the perfect information game. Our induction hypothesis consists of three parts (note that when we talk about \bar{G}' and \underline{G}' in the statement of the induction hypothesis, we identify G' with its corresponding normal form game):

1. For any subgame G' , if \bar{G}' exists, its value is $\bar{v}_{G'}$, where $\bar{v}_{G'}$ is the optimistic value for the corresponding node in the tree. Otherwise, $\bar{v}_{G'} = v_{G'}$.

2. For any subgame G' , if \underline{G}' exists, its value is $\underline{v}_{G'}$, where $\underline{v}_{G'}$ is the pessimistic value for the corresponding node in the tree. Otherwise, $\underline{v}_{G'} = v_{G'}$.
3. The conditions described in Theorem 4 are *necessary* for being an induced normal form proper equilibrium of the game.

After establishing the validity of the statements by induction, we separately prove that the conditions of Theorem 4 are also sufficient for being an induced normal form proper equilibrium.

The statements of the induction hypothesis are true for the trivial subgames of the leaves so we only need to describe the induction step. That is, we are given a game G with r immediate successors G_1, \dots, G_r to the root defining subgames. We assume that the induction hypothesis holds true for the subgames G_1, G_2, \dots, G_r and establish that it holds for G . The first and second statement of the induction hypothesis implies that the optimistic and pessimistic values $\bar{v}_{G'}$ and $\underline{v}_{G'}$ are as as desired for all subgames G' of G except possibly the game G itself. Also, given an induced normal form proper equilibrium for G , Theorem 19 implies that the induced equilibria on all subgames are also induced normal form proper. Then, the third statement of the induction hypothesis implies that the conditions of Theorem 4 are necessary for all nodes except possibly the root. This means that to reestablish the induction hypothesis, we just need to relate \underline{v}_G to the value of \underline{G} , relate \bar{v}_G to the value of \bar{G} , and establish that the behavior in the root node must be as stated in Theorem 4 in any induced normal form proper equilibrium. Also, note that the first condition of Theorem 4 is the property of being a subgame perfect equilibrium. As an induced normal form proper equilibrium is also subgame perfect (by Theorem 19), we only need to show the other two conditions as being necessary when doing the induction.

The induction step is a case analysis by the type (Player 1, Player 2, or Chance) of the root node of the game. For $i = 1, \dots, r$, let $v_i, \bar{v}_i, \underline{v}_i$ be the value, optimistic value and pessimistic value of the subgame G_i .

First assume that the root belongs to Chance. In this case, there is no behavior in the root node, so we only have to reestablish the first and second bullet of the induction hypothesis. We establish the second bullet; establishing the first is a completely symmetric argument. To identify \bar{G} , we determine the superfluous strategies of Player 2. These are exactly the union of superfluous strategies in each of the subgames G_i (assuming, without loss of generality, that the chance node chooses each subgame G_i with strictly positive probability α_i). If there are no such strategies, we have by induction that $\bar{v}_i = v_i$ for all v_i . Note that in this case we have immediately reestablished the second bullet of the induction hypothesis: \bar{G} does not exist and $\bar{v}_G = v_G$. Otherwise, we have to compute the value of \bar{G} . To play minimax in \bar{G} , Player 2 must choose in which subgame G_j to play a superfluous strategy and then play a minimax strategy for \bar{G}_j in that game and a minimax strategy for G_i in any other subgame G_i . The resulting payoff to Player 1 in equilibrium is $\alpha_j \bar{v}_j + \sum_{i \neq j} \alpha_i v_i$. Picking j so as to minimize the loss of Player 2, we obtain that the value of \bar{G} is $v_G + \min_{j, \bar{v}_j > v_j} \alpha_j (\bar{v}_j - v_j)$, reestablishing the second bullet of the induction hypothesis.

Next, suppose that the root of the game belongs to Player 1. It is easy to see that the set of Nash equilibria of G are those strategy profiles where Player 1 chooses a subgame G_i for which $v_{G_i} = v_G$ according to some arbitrary probability distribution on these subgames and an equilibrium in G_i is then played. This means that the set of superfluous strategies of Player 1 is exactly the union of superfluous strategies in those subgames G_i with $v_i = v$ with *all* strategies in subgames G_i for which $v_i < v$. Also, the superfluous strategies of Player 2 are those strategies (y_1, y_2, \dots, y_r) where at least *one* y_i is superfluous in G_i .

If there are no such superfluous strategies for Player 1, we have by induction that

$$V_G = (\cup_{j=1}^r \{v_j, \underline{v}_j\}) \setminus \{v_i\} = \emptyset.$$

In this case \underline{v}_G is defined to be v_G , reestablishing the first bullet of the induction hypothesis. If there are superfluous strategies for Player 1, the game \underline{G} is the game where Player 1 must choose a superfluous strategy and Player 2 can choose any mixed minimax strategy. It is easy to see that when playing minimax in \underline{G} , Player 1 attains the value $\max(V_G)$ either by playing a maximin mixed strategy of G_i where $v_i < v$ or by playing a maximin mixed strategy of \underline{G}_i in a subgame where superfluous strategies exists. That is, we have shown that the value of \underline{G} is $\max(V_G) = \underline{v}_G$, as desired.

If there are no superfluous strategies for Player 2 in G , we have by induction that $I_G = \{j \in \{1, \dots, r\} | v_j = v_i \wedge \bar{v}_j > v_j\} = \emptyset$. In this case \underline{v}_G is defined to be v_G , reestablishing the second bullet of the induction hypothesis. If there are superfluous strategies for Player 2, $I_G \neq \emptyset$. Assume for convenience that $I_G = \{1, \dots, r'\}$ for some $r' \leq r$. The game \bar{G} is the game where Player 1 can choose any mixed maximin strategy of G and Player 2 must choose a superfluous strategy. It is easy to see that to play maximin in \bar{G} , Player 1 must choose to play one of the subgames G_i with $i \in \{1, \dots, r'\}$ and then play a maximin strategy of \bar{G}_i . Also, to play minimax, Player 2 must choose in which subgame G_j to choose a superfluous strategy and then select a minimax mixed strategy of \bar{G}_j in that subgame and a minimax mixed strategy of G_i in any other subgame i . These observations imply that to solve \bar{G} , we can conveniently replace play of a subgame with immediately achieving the value of that subgame. That is, we can solve \bar{G} by solving the following matrix game.

$$\begin{bmatrix} \bar{v}_1 & v_2 & v_3 & \cdots & v_{r'} \\ v_1 & \bar{v}_2 & v_3 & \cdots & v_{r'} \\ v_1 & v_2 & \bar{v}_3 & \cdots & v_{r'} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ v_1 & v_2 & v_3 & \cdots & \bar{v}_{r'} \end{bmatrix}$$

where we in fact have that $v_1 = v_2 = \dots = v_{r'} = v_G$. This game has value $v_G + \frac{1}{\sum_{j \in I_G} (\bar{v}_j - v_j)^{-1}}$, and we have reestablished the second bullet of the induction hypothesis. Also, we find that in the optimal solution to the above matrix game, Player 1 must choose row j with probability $\frac{(\bar{v}_j - v_j)^{-1}}{\sum_{j \in I_G} (\bar{v}_j - v_j)^{-1}}$. Applying Theorem 18, this is also the behavior probability by which Player 1 must choose the subgame G_j in any normal form proper equilibrium (and hence in any induced normal form proper equilibrium), thus establishing the third bullet of the induction hypothesis.

Finally, the case of the root belonging to Player 2 is completely symmetric to the case of the root belonging to Player 1 and is dealt with in a similar way.

Having established by induction that the conditions of Theorem 4 are *necessary* for being an induced normal form proper equilibrium, we now need to prove that they are also *sufficient*. We know that at least one induced proper equilibrium π^* in behavior strategies exists and we just established that it satisfies the conditions of the theorem. So consider two different behavior strategy profiles $\pi^*, \bar{\pi}$ both satisfying the conditions. The crucial observation is that the conditions imply that any two such profiles are *payoff equivalent at any subgame*. That is, if the strategy of one of the players is replaced by the same alternative strategy in both π^* and $\bar{\pi}$ and the game is started in any subgame, the resulting expected payoff is the same in the two profiles. Now, a ‘‘hybrid’’ argument similar to the one used in the proof of Lemma 17 implies that if π^* is an induced normal form proper equilibrium, then so is $\bar{\pi}$. This completes the proof of Theorem 4.