Infinite Coordination Games

Dietmar Berwanger

Laboratoire Spécification et Vérification
CNRS & ENS Cachan
94235 Cachan, France
dwb@lsv.ens-cachan.fr

Abstract

We investigate the prescriptive power of sequential iterated admissibility in coordination games of the Gale-Stewart style, i.e., perfect-information games of infinite duration with only two payoffs. We show that, on this kind of games, the procedure of eliminating weakly dominated strategies is independent of the elimination order and that, under maximal simultaneous elimination, the procedure converges after at most ω many stages.

1 Introduction

Modern computing systems should interact successfully with the environment and never break. As a natural model for non-terminating interactive computation, extensive games of infinite duration have proved to be a suitable analytic framework. For such games, a vast and effective theory has been developed over the past fifty years at the intersection between logic and game theory (for a survey, see [9]). The fundamental model at the basis of this development are Gale-Stewart games [6]: perfect-information games between two strictly competing players with two possible payoff values: win or lose. This basic model has been successfully extended into various directions, including multi-valued payoffs, stochastic effects, partial information, player aggregation, etc. As a common feature most of these extensions postulate a strictly competitive setting.

One major challenge for the analysis of interactive systems consists in handling multiple components that are designed and controlled independently. One can interpret the transition structure of such a system as a game form for several players, each identified with a component, and derive the utility function of each player from the specification of the corresponding component. Via this interpretation, rational strategies in the game correspond to sound designs for components. However, this translation gives rise to infinite non-zero sum games, the theory of which is yet in an initial phase of development. (See [10], for a recent study on Nash equilibrium and refinements in this framework.)

Taking a point of view diametrically opposed to pure conflict models, we investigate extensive games of infinite duration where all participating players receive a common payoff. The players, there may be two or more, thus aim at coordinating their behaviour towards achieving a mutually beneficial outcome. For our analysis, we preserve the remaining aspects of the Gale-Stewart model and restrict our attention to infinite coordination games of perfect information with only two possible payoffs.

Our focus on coordination is motivated by a recurring pattern in the analysis of open systems, in which several components are conceived as a team acting against an adverse environment [1, 16, 12]. Traditionally, such systems are modelled as two-player zero-sum games, and the problem is to construct a strategy for each team member so that the interplay of these distributed strategies guarantees an optimal outcome against the environment. In general, however, the profile of distributed strategies is synthesised by a centralised instance, the designer of the open system, who effectively acts as an external coordinator.

As a far-range objective, we aim at developing an alternative approach to synthesising interaction within a team of players, where the members are themselves responsible for constructing optimal strategies, without involving an external coordinator. Here is the motivating scenario for our investigation. To build a multi-component system, the system designer distributes to different agents a game form representing the possible transitions within a system, and a utility function specifying the desired behaviour of the global system. Each agent is in charge for one component. Independently of the other agents, he should provide an implementation that restricts the behaviour of this particular component in such a way that the composed system satisfies the specification. It is common knowledge among the agents that they all seek to fulfill the same specification, but they are not able to communicate on implementation details, nor to rely on the way in which the game model is represented; this is because they may have different input formats which allow them to reconstruct the original model only up to isomorphism. To accomplish their task, the agents obviously need to share some basic principle of rationality. Our aim is to find principles that are possibly simple and efficient.

In game-theoretic terms, proposing a procedure for resolving this problem amounts to defining a solution concept for coordination games. The concept should prescribe, individually to each player, a set of strategies. Hence, the global solution should be a *rectangular* set: any profile composed of strategies that appear as part of a solution should also constitute a solution.

On finite game trees, coordination games with perfect information and binary payoffs are disconcertingly simple. They can be solved by backwards induction yielding subgame-perfect equilibria, all of which are Pareto efficient, i.e., they attain the maximum available payoff. An equivalent solution is obtained through iterated elimination of weakly dominated strategies.

In the infinite setting, it is a-priori less clear which solution concept would be appropriate. Subgame-perfect equilibria always exist, but they may not form a rectangular set, and prescribing the players to choose a subgame-perfect equilibrium independently could thus lead to coordination failure. The binary payoff scheme induces wide-ranging indifference among the outcomes, offering no grip to refinements based on payoff perturbations. For the same reason, forward-induction arguments do not apply either.

We analyse iterated admissibility, i.e., elimination of weakly dominated strategies, as a solution concept for infinite coordination games. The procedure has been shown to be sound for infinite perfect-information games with two payoffs [4]. Here, we consider a sequential variant of admissibility and show that, on coordination games, it enjoys two desirable properties, that do not hold in the general case.

(i) For any game, the procedure of maximal elimination of dominated strategies converges in at most ω many stages to a non-empty set.

(ii) The outcome of the procedure does not depend on the order of elimination (up to renaming of strategies and deletion of duplicates).

Besides constituting a meta-theoretical criterion for the stability of the proposed solution, order independence is crucial for our application area. If the solution was sensitive to the elimination order, the system designer would need to optimise over different orders, which is a very difficult task.

Applying the procedure towards solving infinite coordination games, we prove, on the positive side, that games with an essentially winning subgame are solvable, i.e., iterated admissibility delivers a rectangular set of strategies, the combination of which always yields the maximal payoff. On the negative side, we show that this classification is tight: if no player has a winning strategy that does not involve the cooperation of other players, admissibility cannot avoid coordination failure.

Our proof is based on a potential characterisation of coordination games. This characterisation also implies that, on infinite coordination with binary payoffs, iterated admissibility provides a refinement of subgame-perfect equilibrium which favours secure equilibria, where a player's payoff cannot decrease under any deviation of other players.

To justify the restrictions assumed for our present model, we point out that the most straightforward relaxations lead to complications that raise doubts on whether admissibility can serve as a meaningful solution concept for more general classes of infinite games. Nevertheless, the question whether the good properties of infinite coordination games with two payoffs can be extended to games with finitely many payoffs remains open. We show that, unlike the case of finite coordination games with perfect information, or infinite non-zero games with two payoffs, already a few payoffs are sufficient to generate forward-induction effects in infinite coordination games, which appear to take the analysis out of the reach of our present methods.

2 Formalities

In situations that involve n players, we refer to a list of elements $x = (x^i)_{i<n}$, one for each player, as a *profile*. For any such profile we write x^{-i} to denote the list $(x^j)_{j<n, j \neq i}$ of elements in x for each player except i. Given an element x^i and a list x^{-i}, we denote by (x^i, x^{-i}) the profile $(x^i)_{i<n}$. For clarity, we will always use superscripts to specify to which player an element belongs. If not quantified otherwise, we usually refer to player i meaning *any* player.

2.1 Coordination games

An extensive *coordination game* for n players is a structure $\Gamma = \big(\mathcal{T}, u\big)$, where $\mathcal{T}$ is a directed tree over a domain T of elements called *positions*, equipped with a partition $(T^i)_{i<n}$ of its non-terminal nodes, and u is a *utility* function that maps every maximal path through $\mathcal{T}$ to an integer number.

Game trees may be of arbitrary branching and of depth up to ω. We write $\preceq$ for the partial order associated to $\mathcal{T}$. To play the game, the players form a maximal path through $\mathcal{T}$ starting from the root. Whenever a non-terminal position p is reached, the player i with $p \in T^i$ is in turn to prolong the path by choosing one of the outgoing edges. The outcome of a play is thus a maximal path π through the game tree, and its utility $u(\pi)$ determines the common payoff that all players receive. We identify *plays* with maximal paths through $\mathcal{T}$; an *initial* play is a prefix of a play.

A *strategy* for player i is a function $s^i : T^i \to T$ that associates to every position in T^i a successor in $\mathcal{T}$. We denote the set of all strategies for player i by S^i, and the set of all strategy profiles $\times_{i<n} S^i$ by S. An (initial) play $p_0, p_1, \dots$ *follows* a strategy $s^i \in S^i$, if for every $p_k \in T^i$, we have $p_{k+1} = s^i(p_k)$. Any strategy profile $s = (s^i)_{i<n}$ determines a unique play π which follows all of its components; we often write $u(s)$ to mean $u(\pi)$.

A coordination game is represented in *normal form* by a profile $S = (S^i)_{i<n}$ of strategy sets together with a utility function $u : S \to \mathbb{Z}$. Given a game $\Gamma = (S, u)$ in normal form, and a profile Q of non-empty sets $Q^i \in S^i$, for all $i < n$, the *restriction* of Γ to Q is the normal-form game over the strategy sets $(Q^i)_{i<n}$ with the utility function u restricted to Q. When Γ is fixed, we identify the restriction (Q, u) of Γ with its strategy space Q.

Given an extensive game Γ and a position p, the *subgame* Γ_p is the extensive game obtained by restricting all components of Γ to positions $p' \succeq p$. We denote the restriction of a strategy $s^i \in S^i$ to this domain by $s^i|_p$ and we write $u_p(s)$ for the utility of a profile $s|_p := (s^i|_p)_{i<n}$ in Γ_p.

Games with binary payoffs. Our basic model are games with only two possible payoffs: 1 (win) and -1 (lose). For such a game $\Gamma = (S, u)$, we say that a play π is *winning* if $u(\pi) = 1$. Likewise, a profile $s \in S$ is winning (for all players) if the play determined by s is winning. When talking about a restriction $Q \subseteq S$ of Γ, we say that a strategy $s^i \in S^i$ is *winning with respect to* Q, if any profile (s^i, t^{-i}) with $t^{-i} \in Q^{-i}$ is winning. Likewise, we say that s^i is *losing* with respect to Q, if no profile (s^i, t^{-i}) with $t^{-i} \in Q^{-i}$ is winning.

2.2 Sequential Admissibility.

We define a sequential version of iterated admissibility based on simultaneous maximal elimination of weakly dominated strategies, following Bicchieri and Schulte [5].

Given two strategies $s^i, r^i \in S^i$ and a set $Q^{-i} \subseteq S^{-i}$, we say that s^i *dominates* r^i on Q^{-i} at a position $p \in T^i$, if

$$u_p(s^i, t^{-i}) \geq u_p(r^i, t^{-i}) \text{ for all } t^{-i} \in Q^{-i}, \text{ and}$$
$$u_p(s^i, t^{-i}) > u_p(r^i, t^{-i}) \text{ for some } t^{-i} \in Q^{-i}.^{1}$$

For a set $Q \subseteq S$ and a position $p \in T^i$, we say that a strategy $s^i \in S^i$ is *admissible* on Q at p, if no strategy in Q^i dominates s^i on Q^{-i} at p. A strategy $s^i \in S^i$ is *sequentially admissible* on Q, if it is admissible on Q at every position $p \in T^i$.

For a game Γ, we define simultaneously for all players i:

- $Q^i_0 := S^i$;
- $Q^i_{\alpha+1} := \{s^i \in Q^i_\alpha : s^i \text{ is admissible on } Q_\alpha\}$, for every successor ordinal $\alpha + 1$, and
- $Q^i_\lambda := \bigcap_{\alpha<\lambda} Q^i_\alpha$, for every limit ordinal λ.

A strategy $s^i \in Q^i_\alpha$ is called α-*admissible*. As the stages are decreasing, the iteration reaches a fixed point $(Q^i_\infty)_{i<n}$. A strategy $s^i \in Q^i_\infty$ is called *iteratively admissible*. Observe that, at each stage, the set Q_α of admissible strategies yields a rectangular set, i.e., a Cartesian product of sets.

[1] Whenever we refer to dominance throughout the paper, we actually mean *weak* dominance, in the classical sense. There is no risk of confusion, as the notion of strict dominance plays no role when we are concerned about pure strategies in games with only two payoffs.

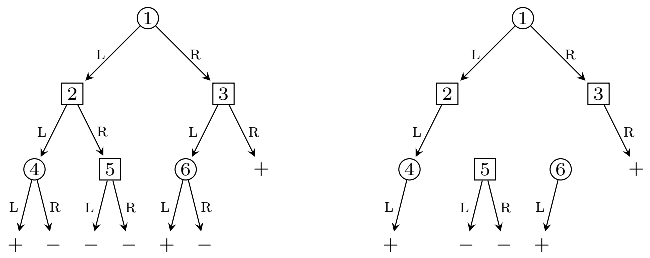

(a) Game tree (b) Admissible strategies

Figure 1: Finite coordination game

3 Solvable and unsolvable games

In the following two sections, we restrict our attention to infinite coordination games with only two payoffs.

A coordination game Γ is *solvable*, if the utility function is constant on Q^i_∞, that is, if all iteratively admissible profiles yield the same payoff. In any solvable game where the winning set is nonempty, iterated admissibility thus guarantees a winning outcome. In contrast to this, if a game is not solvable then, by the definition of admissibility, for every iteratively admissible strategy of a player there exist iteratively admissible strategies of other players such that the resulting profile is losing; hence, in this case, iterated admissibility cannot avoid coordination failure.

3.1 Examples

We begin with some examples to illustrate how the elimination procedure works. In the graphical representations throughout the paper, there will be only two players. We draw the positions of Player 0 in circles and those of Player 1 in squares; binary payoffs in $\{-1, 1\}$ are represented by their sign.

The case of finite games is very simple in our setting. In the game of Figure 1, for instance, any strategy of Player 0 which at position 6 chooses R is sequentially dominated by the one which instead chooses L. Likewise, on the full set of strategies of Player 0, any strategy of Player 1 which at 3 chooses L is dominated by the one that instead chooses R. However, the two strategies are incomparable on a set from which Player 1 has previously eliminated the strategies choosing R at 6. In general, for finite games, one round of simultaneous elimination of dominated strategies yields a restriction where losing outcomes are reachable only from positions from which there are no reachable winning ones. Moreover if, for some subgame, a player i has a strategy that guarantees a win regardless of how the other players move, (as it is the case at position 3 in the example), then the admissibility procedure eliminates any strategy of i that is not winning for this

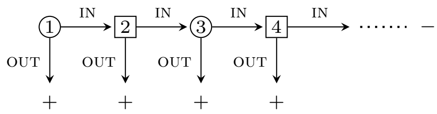

Figure 2: Game of infinite duration

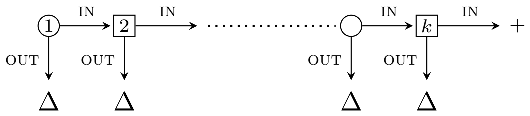

Figure 3: Game Δ_k requiring k iterations

subgame. As a consequence, every finite game is solvable, and the number of rounds needed until the iteration stabilises is bounded by the depth of the game tree.

Figure 2 shows a game with one infinite play which is losing, whereas all finite plays are winning. In the first round, the elimination procedure removes all strategies that persist in playing IN from some position onwards, which readily solves the game and ensures coordination.

However, unlike the finite case, there are infinite games which are not solvable. Consider, for instance the two-player game on the infinite binary tree, where the players strictly alternate in choosing either the left or the right successor of the current position, and the winning condition requires that the sequence of choices be finally constant. In this game, let us call it Δ, no player has a winning strategy. The first round of elimination removes all strategies where a player would keep shifting between left and right, regardless of how his partner moves. However, in the obtained restriction, the players can still not guarantee a win, and none of the surviving strategies are dominated. Thus, the iteration terminates leaving the game unsolved.

We can use the game Δ to build a game Δ_k that requires k iterations as shown in Figure 3. Clearly, the only iteratively admissible strategy in Δ_k is the one that always chooses IN. However, the strategies that choose OUT at position ℓ can be eliminated only at the stage at which all strategies that choose OUT at some later position $\ell' > \ell$ are eliminated. To be solved, the game thus needs k rounds to propagate the winning payoff. Figure 4 extends this example to obtain a game that requires ω many iterations.

3.2 Characterisation

Before we proceed to a formal discussion, let us describe the procedure step-by-step for an arbitrary game Γ with a non-empty winning condition. In

Figure 4: Game Δ_ω requiring ω iterations

the first stage, there may exist subgames Γ_p in which a player has a winning strategy (with respect to S). We call such strategies *strongly winning* in Γ_p. For any subgame Γ_p in which a player i has a strongly winning strategy, all strategies of S^i that are not winning in Γ_p are eliminated. All strategies that are losing (with respect to S) are also eliminated. This yields Q_1.

Notice that no strategy in Q_1 can be losing with respect to Q_1. However, there may exist subgames Γ_p in which a player has a winning strategy with respect to Q_1. We call such strategies *weakly winning.*

In the second stage, if no weakly winning strategy exists, the procedure terminates, concluding that the game is not solvable. Otherwise, for any subgame Γ_p where a player i has a weakly winning strategy, all strategies of Q^i that are not winning in Γ_p with respect to Q_1 are eliminated, thus yielding Q_2. It may happen that no strategies are eliminated in this stage (note that strongly winning strategies are also weakly winning). Then, the procedure terminates, concluding that the game is solvable.

In the following stages $\alpha \geq 2$, strategies are eliminated only due to the discovery of new subgames where a player has a winning strategy with respect to the previous stage. These subgames arise by backwards propagation of winning subgames discovered in the first two stages. The propagation continues until it reaches the root of the initial game or of a previously solved subgame. Every subgame where a player has a (weakly) winning strategy will thus propagate for a finite number of stages. Since winning subgames may occur arbitrarily deep in the tree, the number of stages to propagate them cannot be bounded. Therefore, the iteration may require ω many steps to stabilise.

Based on the insights from the above description, we advance the following statement which will be proved in the remaining part of the section.

Theorem 3.1 (Constructiveness)**.** For any coordination game with binary payoffs, the procedure of simultaneous maximal elimination of dominated strategies converges in at most ω many steps. In case the game is not solvable, the procedure terminates at the second stage.

It becomes apparent that, for coordination games, the elimination procedure is considerably simpler than in the case when payoffs are not common. Indeed, losing strategies can occur only at the initial stage. Accordingly, each subgame that contains no subgames where a player has a weakly winning strategy freezes from the second stage onwards: every strategy of a player in such a subgame can lead to a winning or a losing outcome, depending on the choice of the other players. None of the players is able to avoid this interdependence and, therefore, no strategy can be eliminated. This leads us to the following characterisation of solvable games as being, essentially, those games where coordination effort is required only for finitely many steps at the beginning of a play.

Theorem 3.2 (Solvability)**.** A game Γ with binary payoffs is solvable if, and only if, it contains a subgame where one of the players has a weakly winning strategy, i.e., a strategy that is winning with respect to the set of 1-admissible strategies.

3.3 Proof

Our justification of Theorems 3.1 and 3.2, and the analysis of order dependency in the following section, relies on the notion of potential of (a restriction of) a coordination game. Intuitively, the potential of a restricted game describes the payoff that the players can achieve when confining themselves to playing within the restriction

The notion of potential is based on the value function defined in [4] for characterising the dominance relation. Let us fix an extensive coordination game Γ for the rest of the section. The *value* of a restriction $Q \subseteq S$ of Γ for player i is a function χ^i that assigns to every position $p \in T$ a value $\chi^i_p(Q) \in \{-1, 0, 1\}$, as follows:

- $\chi^i_p(Q) = 1$ (winning) if there exists a strategy $s^i \in Q^i$ such that $u_p(s^i, t^{-i}) = 1$, for all $t^{-i} \in Q^{-i}$,
- $\chi^i_p(Q) = -1$ (losing) if $u_p(s) = -1$, for every profile $s \in Q$, and
- $\chi^i_p(Q) = 0$ (undetermined), otherwise.

The following lemma is an adaptation of the Value Characterisation from [4, Lemmata 8 and 9] for the sequential variant of dominance.

Lemma 3.3 (Value characterisation [4])**.** Let $Q \subseteq S$ be an iteration stage of the admissibility procedure.

(i) A strategy $s^i \in S^i$ dominates $r^i \in S^i$ on Q if, and only if,

  – $\chi^i(s^i, Q^{-i}) \geq \chi^i(r^i, Q^{-i})$, and
  – $\chi^i_p(s^i, Q^{-i}) > \chi^i_p(r^i, Q^{-i})$, for some $p \in T$.

(ii) A strategy $s^i \in S^i$ is sequentially admissible on Q if, and only if, $\chi^i(s^i, Q^{-i}) \geq \chi^i(Q)$.

In the special case of coordination games, the value functions of the different players are intimately connected. Obviously, if $\chi_p^i(Q) = -1$ for one player i, then $\chi_p^j(Q) = -1$ for all players j. Moreover, if $\chi_p^i(Q) = 1$ for a player i, then for the set Q' obtained from Q by maximal simultaneous elimination of dominated strategies, we have $\chi_p^j(Q') = 1$ for all players j.

We define the *potential* $\chi_p(Q)$ of a restriction $Q \subseteq S$ at position p to be maximum value $\chi_p^i(Q)$ over all players $i < n$. The potential function induces a partial order over restrictions of Γ: for $Q, Q' \subseteq S$, we write $\chi(Q) \leq \chi(Q')$ if $\chi_p(Q) \leq \chi_p(Q')$, for all $p \in T$. Likewise, we denote by $\max(\chi(Q), \chi(Q'))$ the function that assigns to every $p \in T$ the value $\max\{\chi_p(Q), \chi_p(Q')\}$. When one of the sets Q^i is a singleton $\{s^i\}$ we write $\chi(s^i, Q^{-i})$ rather than $\chi(\{s^i\}, Q^{-i})$.

According to Lemma 3.3, the restrictions Q where all strategies are admissible on Q itself are characterised by $\chi(s^i, Q^{-i}) = \chi(Q)$ for all players $i < n$ and all $s^i \in Q^i$. The fixed points of the elimination procedure are examples of such restrictions.

In contrast to the general case, in which the value function may increase and decrease along the elimination stages, the value – and thus the potential – of coordination games is monotonic. Indeed, the potential observed at a particular position p during the iteration can change only from 0 (undetermined) to 1 (winning).

Lemma 3.4 (Monotonicity, stationary values)**.** For any two iteration stages $Q \supseteq Q' \supsetneq \varnothing$ of the admissibility procedure, the following properties hold:

(i) $\chi(Q) \leq \chi(Q')$, and

(ii) if $\chi_p(Q) \neq 0$ then $\chi_p(Q') = \chi_p(Q)$, for any position $p \in T$.

Proof. The statement is a consequence of the Lemma 3.3. Here, we will only sketch the argument. In case $\chi_p(Q) = -1$, every strategy profile in $Q' \subseteq Q$ is losing, hence $\chi_p(Q') = -1$. In case $\chi_p(Q) = 1$, there exists a strategy $s^i \in Q^i$ which is winning in the subgame Γ_p restricted to Q. Then, any later stage Q' contains a strategy r^i that agrees with s^i on Γ_p, hence $\chi_p(Q') = 1$. Finally, in case $\chi_p(Q) = 0$, there exists a profile $s \in Q$ with $u_p(s) = 1$. To derive a contradiction, suppose that $\chi_p(Q) = -1$. Then, there exists a player i for which all strategies in Q^i that agree with s^i on the play π determined by s in Γ_p are eliminated. This can happen only if, at some stage Q'' between Q and Q', player i has a winning strategy r^i for a subgame $\Gamma_{p'}$ rooted at a position p' on π; let us choose i and p' such that p' is of minimal distance from p. Then, any play in Q'' that proceeds along π up to p' and then follows the strategy r^i is winning, witnessing that $\chi_p(Q'') \geq 0$. When reiterating the argument, the profitable deviations propagate towards the root, and we obtain $\chi_p(Q') = 1$, a contradiction. Q.E.D.

The potential function satisfies the following consistency property along the edges of the game tree.

Lemma 3.5 (Propagation)**.** Let $Q \supseteq Q'$ be two consecutive iteration stages of the admissibility procedure on Γ, and let $p \in T$. For p' ranging over the set of direct successors of p in $\mathcal{T}$, we have:

$$\chi_p(Q) \leq \max_{p'} \chi_{p'}(Q) \quad \text{and} \quad \chi_p(Q') \geq \min_{p'} \chi_{p'}(Q).$$

Proof. For an arbitrary player i, the value of Q at p is bounded by

$$\min_{p'} \chi^i_{p'}(Q) \leq \chi^i_p(Q) \leq \max_{p'} \chi^i_{p'}(Q)$$

On the one hand, as the potential is the maximum value, this implies that $\chi_p(Q) \leq \max_{p'} \chi_{p'}(Q)$. On the other hand, $\chi^j_{p'}(Q') \geq \chi^i_{p'}(Q)$ for all players j and all positions p' (eliminating dominated strategies from Q^i is profitable for all players). This implies $\chi_{p'}(Q') \geq \chi^i_{p'}(Q)$ and, as i was chosen arbitrarily, we can conclude $\chi_{p'}(Q') \geq \min_{p'} \chi_{p'}(Q)$. Q.E.D.

In particular, within a restriction Q, we have $\chi_p(Q) = -1$ if, and only if, $\chi_{p'}(Q) = -1$ for all $p' \succeq p$. Also, it follows that, for any position p, either p is a leaf with $\chi_p(Q)$ corresponding to its payoff, or there exists a direct successor p' with $\chi_{p'}(Q) \geq \chi_p(Q)$.

From Lemmas 3.4 and 3.5, we can immediately derive the following properties of individual elimination steps.

Lemma 3.6 (Potential transformation)**.** For any two consecutive iteration stages $Q \supseteq Q'$ of the admissibility procedure, the following properties hold:

(i) $\chi(Q) = \chi(Q')$, if, and only if, all strategies in Q' are admissible on Q'.

(ii) For any position p, if $\chi_{p'}(Q) \leq 0$, for all $p' \succeq p$, then $\chi_{p'}(Q') = \chi_{p'}(Q)$ for all $p' \succeq p$.

(iii) If $\chi_p(Q) = 0$ and $\chi_{p'}(Q) = 1$ for some $p' \succeq p$, then there exists a position p'' with $p \preceq p'' \prec p'$ such that $\chi_{p''}(Q) = 0$ and $\chi_{p''}(Q') = 1$.

On the basis of these properties we can describe, in terms of potential, how the strategy elimination proceeds and how it terminates. First, if no winning strategy for a subgame Γ_p is encountered at stages Q_0 or Q_1, that is, if $\chi_p(Q_1) \leq 0$ for all p, then the iteration terminates with Q_1, by Lemma 3.6 (i) and (ii). Otherwise, consider the upwards closure $X := \{p' \in T : \chi_p(Q_1) = 1 \text{ for some } p \succ p'\}$ of the positions that are winning in Q_1. On the one hand, for all positions $p \in T \setminus X$, the potential χ_p is stationary from stage Q_1 onwards, by Lemma 3.6 (ii). On the other hand, Lemma 3.6 (iii) implies that every position $p \in X$ will have $\chi_p(Q_\alpha) = 1$ for some finite ordinal α (winning positions travel upwards). Hence, after at most ω many steps $\chi(Q_\alpha)$ will become stationary and the iteration terminates with $\chi_p(Q_\infty) = 1$ if, and only if, $p \in X$. This concludes the proof of Theorems 3.1 and 3.2.

4 Order independence

One way to argue why iterated admissibility is a meaningful solution concept is by illustrating that the players may identify dominated strategies from previous or fictive plays of a game and learn to avoid them henceforth. By taking different parts in the game, they would also understand that other players will avoid dominated strategies, which leads to iterating the elimination procedure until the evolution reaches a fixed point [14].

This view, however, raises the criticism that the solution concept would depend on the order in which strategies are eliminated. The choice of eliminating all dominated strategies at once for all players may seem arbitrary. Indeed, there are simple examples of games with payoffs that are not common to the players, where a procedure in which the players take turns while eliminating strategies would lead to a different outcome [18].

Over the past thirty years, several conditions to characterise finite extensive games that are immune against this criticism have been proposed [11, 17, 7, 15]. Typically, these conditions feature a form of *transference of decisionmaker indifference* requiring that no player should be indifferent between choices which affect other player's payoffs. In the kind of games we consider here, this condition obviously holds. Nevertheless, there are coordination games with infinitely many payoffs, where the outcome of iterated elimination of dominated strategies depends of the elimination order, as we shall see in 5. It turns out that, for games with binary payoffs, the order in which dominated strategies are eliminated does not matter when solving infinite coordination games via iterated admissibility.

To make this statement more precise, we introduce the necessary notions following Marx and Swinkels [15].

A *reduction sequence* for a game Γ is a descending sequence $(R_\alpha)_{\alpha \in \mathrm{On}}$ of rectangular sets of strategy profiles satisfying the following conditions:

- $R_0 := S$;
- $R_{\alpha+1} \subseteq R_\alpha$ for every successor ordinal $\alpha + 1$, and each $s^i \in R^i_{\alpha+1} \setminus R^i_\alpha$ is dominated on R_α, and
- $R_\lambda := \bigcap_{\alpha<\lambda} R_\alpha$, for every limit ordinal λ.

We denote the fixed point of the sequence with R_∞. A reduction sequence is *full*, if no strategy in R_∞ is dominated.

Notice that the stages of maximal simultaneous elimination of dominated strategies used to define admissibility form a particular (full) reduction sequence. However, arbitrary reduction sequences may decrease over uncountably many stages before reaching R_∞. Consider, for instance, a 2-player game on the infinite binary tree with only one winning play. Then, the full reduction sequence which removes the dominated strategies one by one is of length at least ω_1.

We remark here, that the Value-Characterisation Lemma 3.3 and the Monotonicity Lemma 3.4 still hold when we consider stages of arbitrary reduction sequences, instead of stages of simultaneous maximal elimination. This can be shown following the corresponding proofs in [4].

Any full reduction sequence of a game Γ induces a restriction of the game to the strategy space $(R^i_\infty)_{i<n}$. We call such a restriction a *reduction* of Γ. To compare different reductions of a game, we use a notion of game equivalence that allows renaming of (normal-form) strategies and removal of redundant ones.

Definition 4.1 (Game equivalence)**.** Let R be a restriction of Γ. Two strategies $s^i, r^i \in R^i$ are *payoff equivalent* on R, written as $s \sim_R r$, if $u_p(s^i, t^{-i}) = u_p(r^i, t^{-i})$, for every $t^{-i} \in R^{-i}$ and all positions $p \in T$.

For two restrictions R, R' of Γ, we say that R is *embeddable* into R', if there exists an embedding $f : R \to R'$ which respects payoff equivalence:

(i) $u_p(s_0, \ldots, s_{n-1}) = u_p(f(s_0), \ldots, f(s_{n-1}))$, for all $s \in R$ and $p \in T$;

(ii) $f(s^i) \sim_{R'} f(r^i)$ if, and only if, $s^i \sim_R r^i$, for all $i < n$ and $s^i, r^i \in R^i$.

We say that two restrictions R, R' are *equivalent* if R is embeddable into R' and, vice-versa. In this case, we write $R \equiv R'$.

Clearly, both $\sim$ and $\equiv$ are equivalence relations, the former over strategies and the latter over game restrictions. To recover the sense in which $\equiv$ captures a notion of equivalence up to renaming of strategies and removal of redundant ones, consider two restrictions $R \equiv R'$ and transform each of them by keeping only one strategy from each payoff-equivalence class. The residual games on $\hat{R}$ and $\hat{R}'$ obtained in this way are still equivalent, i.e., there exist injective embeddings between them. By the Cantor-Bernstein-Schröder argument, this implies that there exists an isomorphism between the game structures obtained by restricting Γ to $\hat{R}$ and $\hat{R}'$, respectively.

We remark that, if the players do not receive the same payoff, different reduction sequences may lead to non-equivalent outcomes already in finite extensive games with binary payoffs. Consider, for instance, the game in Figure 5(a) with one non-common payoff value $\pm$ meaning that Player 0 (circle) wins and Player 1 (square) loses. The two reduction sequences obtained by eliminating single strategies in the order described in Figure 5(b)

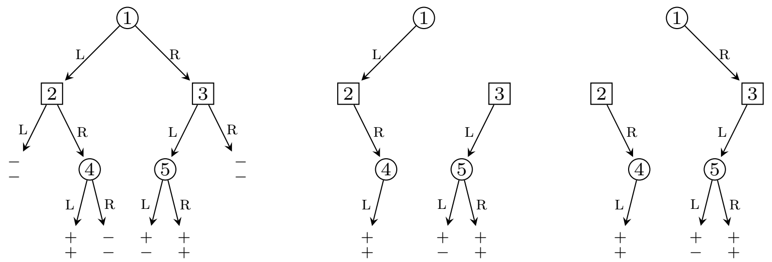

(a) Game tree (b) 4R, 2L, 1R, 3R (c) 3R, 1L, 2L, 4R

Figure 5: Elimination orders with different outcomes

and 5(c) lead to different outcomes for Player 1: in the former reduction he wins all plays, whereas he may lose when RLL is played according to the latter one.

Our aim is to prove that, in coordination games with two payoffs, all procedures of iterative elimination of dominated strategies lead to a unique outcome, up to equivalence.

Theorem 4.2 (Order independence)**.** Any two reductions of an infinite coordination game with binary payoffs are equivalent.

4.1 Proof

Following the unifying approach of Apt [2] for proving order-invariance of finite games, we formulate our proof of Theorem 4.2 in terms of abstract reduction systems. As we deal with strategy spaces of infinite cardinality, we cannot rely on Newman's Lemma which is the basic principle of this framework. Instead, we establish that the reduction sequences which we consider here, satisfy a transfinite confluence property.

For the sequel of this section, let us fix a game Γ. We use the notation $R \twoheadrightarrow R'$ to express that there exists a (possibly transfinite) reduction sequence with $R_0 = R$ and $R_\alpha = R'$ for some $\alpha \in \mathrm{On}$.

First, we introduce a notion of game equivalence which is weaker than $\equiv$, but easier to handle. Given two restrictions of Γ, we define $R \equiv_\chi R'$ to hold if $\chi(R) = \chi(R')$. Clearly, $R \equiv R'$ implies $R \equiv_\chi R'$, but the converse is in general not true. For fixed points of full reduction sequences, however, it turns out that the two notions coincide.

Lemma 4.3. Let $R, R' \subseteq S$ be reductions of Γ. Then $R \equiv R'$ if, and only if, $R \equiv_\chi R'$.

Proof. Let $X \subseteq T$ be the set of positions with $\chi_p(R) = \chi_p(R') = 1$. As all strategies in R are admissible on R, this set is upwards closed with respect to $\preceq$, i.e., if $p \in X$ then $p' \in X$ for all $p' \preceq p$.

To define a suitable embedding f of R into R', consider any strategy $r^i \in R^i$ for a player i, and pick an arbitrary strategy $r'^i \in R'^i$. Let $f(r^i)$ be the strategy that agrees with r^i on X and with r'^i on $T \setminus X$.

We claim that $f(r^i) \sim_S r'^i$, that is $u_p(f(r^i), t^{-i}) = u_p(r'^i, t^{-i})$ for all $t^{-i} \in S^{-i}$. This is sufficient to show that the image $f(r^i)$ belongs to R'^i, since for any restriction $Q \subseteq S$, and in particular for the stages of $S \twoheadrightarrow R'$, the set of strategies that are admissible on Q is closed under $\sim_S$.

To prove the claim note that, for any position $p \in T \setminus X$, the strategies $f(r^i)$ and r'^i agree on $\{p' : p' \succ p\}$, by construction. As $T \setminus X$ is downwards closed, we have $u_p(f(r^i), t^{-i}) = u_p(r'^i, t^{-i})$. For positions in $T \setminus X$, we argue as follows. Since r^i and r'^i are admissible on R and R', respectively, we have, by Lemma 3.3, $\chi(r^i, R^{-i}) = \chi(R^{-i})$ and $\chi(r'^i, R'^{-i}) = \chi(R'^{-i})$. By hypothesis, these potentials are equal. In particular, for any $p \in X$, we have $\chi_p(r^i, R^{-i}) = \chi_p(r'^i, R^{-i}) = 1$. On the other hand, both r^i and r'^i are image-closed on X, i.e., $r^i(p) \in X$ for any $p \in X$ and likewise for r'^i. Therefore, we have for all $p \in X$ and every $t^{-i} \in S^{-i}$,

$$u_p(f(r^i), t^{-i}) = u_p(r'^i, t^{-i}) = 1 = u_p(r^i, t^{-i}).$$

Next, let us verify that f is an embedding, that is, $u_p(s^0, \dots, s^n) = u_p(f(s^0), \dots, f(s^n))$ for all p and all $s \in R$. At any $p \in T \setminus X$, the condition holds because each strategy s^i agrees with its image $f(s^i)$ on $T \setminus X$; at $p \in X$ all profiles in R and R' yield the same payoff 1.

To justify that f respects payoff equivalence, the argument is similar. Consider a pair of strategies $r^i \sim_R s^i$, that is $u_p(r^i, t^{-i}) = u_p(s^i, t^{-i})$ for all $p \in T$ and $t^{-i} \in R^{-i}$. For $p \in T \setminus X$, the strategies r^i and s^i agree with their f-image on $T \setminus X$, and hence $u_p(f(r^i), t^{-i}) = u_p(f(s^i), t^{-i})$, for every $p \in T \setminus X$ and all $t^{-i} \in R^{-i}$. On the other hand, for $p \in X$, we have, $u_p(f(r^i), t^{-i}) = u_p(f(s^i), t^{-i})$, because $f(r^i)$ and $f(s^i)$ are in R'.

Symmetrically, we can construct an embedding of R' into R that preserves payoff equivalence, thus concluding the proof that $R \equiv R'$. Q.E.D.

In the next step, we relate reduction sequences obtained by arbitrary elimination of dominated strategies with the sequence $(Q_\alpha)_{\alpha \in \mathrm{On}}$ obtained by maximal simultaneous elimination of dominated strategies.

We define the *rank* $|p|$ of a position $p \in T$ to be the least ordinal α such that $\chi_p(Q_\alpha) < \chi_p(Q_{\alpha+1})$ if such an ordinal exists, and 0 otherwise. Notice that for all positions $p \in T$ with $\chi_p(S) \leq 0$, we have $|p| = 0$, according to Lemma 3.6(ii). Likewise, any position p such that there exists a strongly or weakly winning strategy in $\Gamma|_p$ has rank 0 or 1, respectively. Essentially, the rank of a position represents the stage at which it is discovered as being winning, losing, or undetermined during the process of maximal simultaneous elimination.

Intuitively, the following lemma points out that any set of winning positions discovered through arbitrary elimination is also discovered through maximal simultaneous elimination.

Lemma 4.4 (Covering). If $S \twoheadrightarrow R$, then $\chi(R) \leq \chi(Q_\kappa)$, for the stage $\kappa = \sup\{|p| : p \in T \text{ with } \chi_p(R) = 1\}$.

Proof. If there are no positions $p \in T$ with $\chi_p(R) = 1$, we simply have $\chi(R) = \chi(S) = \chi(Q_0)$. Otherwise, for any position p with $\chi_p(Q_\infty) = -1$, we have $\chi_p(R) = -1 \leq \chi_p(Q_\kappa)$. For positions p with $\chi_p(Q_\infty) = 1$, we have $\chi_p(Q_0) \geq 0$ and thus, by monotonicity $\chi_p(Q_\kappa) \geq 0$; due to the way k was chosen, if $|p| \leq \kappa$ then $\chi_p(Q_\kappa) = 1$ and, if $|p| > \kappa$, then $\chi_p(R) = 0$. Thus, $\chi_p(Q_\kappa) \geq \chi_p(R)$, in both cases. Finally, for positions p with $\chi_p(Q_\infty) = 0$ we argue by transfinite induction along the reduction sequence, that no undetermined strategy is ever eliminated in the subgame starting at p, that is, $\{s^i|_p : s^i \in R_1^i\} \subseteq \{s^i|_p : s^i \in Q_\alpha^i\}$, for any stage α and every player $i < n$. Hence, we can conclude that $\chi_p(R) = 0 = \chi_p(Q_\kappa)$.

Q.E.D.

The next lemma, illustrated in Figure 6(a), shows that, when we set out on a particular reduction sequence $S \twoheadrightarrow B$, there is no risk to miss a position p which would have been discovered following a different sequence; it is always possible to continue the reduction from B onwards such that p will be discovered in at most ω many steps.

Lemma 4.5 (Catch up). If $B \twoheadleftarrow S \twoheadrightarrow C$, then

(i) there exists an index δ such that $\chi(Q_\delta) \geq \max\{\chi(B), \chi(C)\}$, and

(ii) for the least such δ, there exists $D \twoheadleftarrow B$ such that $\chi(D) = \chi(Q_\delta)$.

Proof. Let $\delta := \sup\{|p| : p \in T \text{ with } \chi_p(B) = 1 \text{ or } \chi_p(C) = 1\}$. According to Lemma 4.4, $\chi(Q_\delta) \geq \max\{\chi(B), \chi(C)\}$ and there is no index smaller than δ with this property.

To prove the second point, let us construct a sequence $(B_\alpha)_{\alpha<\delta}$ with $B_0 := B$ and $B_{\alpha+1}$ consisting of those strategies $s^i \in B_\alpha$ that are admissible on B_α at all positions of rank at most α; for the case $\delta = \omega$, we set $B_\omega := \bigcap_{\alpha<\omega} B_\alpha$.

We claim that the set $D := B_\delta$ has the required properties. Clearly, $D \twoheadleftarrow B$ via a reduction sequence of length $\delta \leq \omega$. For every position p of rank up to δ and every index $\alpha \leq \delta$, by construction, $\chi_p(B_\alpha) = \chi_p(Q_\alpha)$. Thus, for all these positions, $\chi_p(D) = \chi_p(Q_\delta)$. On the other hand, for each position p with rank greater than δ, we have $\chi_p(D) = \chi_p(B) = 0$, again by Lemma 4.4. Hence, overall $\chi(D) = \chi(Q_\delta)$. Q.E.D.

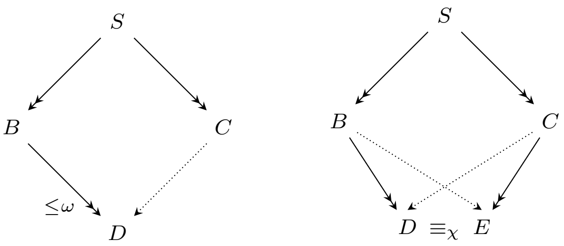

(a) Catch up (b) Confluence

Figure 6: Reduction patterns

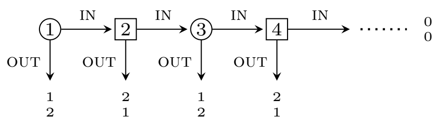

Figure 7: Conflicting interests

Finally, we are ready to state our confluence property for reductions of extensive coordination games.

Lemma 4.6 (Confluence)**.** If $B \twoheadleftarrow S \twoheadrightarrow C$, there exist restrictions $D \twoheadleftarrow B$ and $E \twoheadleftarrow C$ such that $D \equiv_\chi E$.

Proof. It is sufficient to apply the construction of 4.5 on the two reduction paths as illustrated in Figure 6(b). The restrictions D and E obtained in this way then have the same potential as the least stage Q_δ with $\chi(Q_\delta) \geq \max\{\chi(B), \chi(C)\}$. Q.E.D.

By Lemma 4.4, the potential of any reduction is bounded above by $\chi(Q_\infty)$. Hence, for any fixed point R of a full reduction sequence, we have $R \equiv_\chi Q_\infty$. According to Lemma 4.3, it follows that every extensive coordination game has a unique normal form, which concludes our proof of Theorem 4.2.

5 Beyond binary payoffs

From the perspective of the more classical theory of extensive games, the restriction to binary payoffs assumed in Gale-Stewart games may appear too limiting. However, it seems doubtful whether admissibility can yield a sound solution concept for games of infinite duration with a more general payoff structure.

One problem is that, already when there are three different payoffs, one round of maximal simultaneous elimination can remove infinite dominance chains without letting any dominating strategy survive. As an example, consider the game in Figure 7 which represents a coordination situation with conflicting interests. Notice that, on the full set of strategies, any strategy of Player 0 that chooses IN in the first $k-1$ rounds, that is, at positions $1, 3, \dots, 2k-1$, and OUT in round k is dominated by one that keeps choosing IN until round k, and only afterwards, in round $k+1$ chooses OUT. Thus, maximal simultaneous elimination, will eliminate all strategies of Player 0 except the one that chooses IN forever; the same reasoning applies to Player 1. Consequently, the only 1-admissible strategy profile consists in choosing IN forever yielding payoff $(0, 0)$, which is less than any combination of the eliminated strategies would yield for any player. With maximal elimination in sequential order, where the dominated strategies of Player 0 are removed first, Player 1 would retain all strategies that choose OUT after a finite number of rounds, which yields a solution with payoff $(2, 1)$. On the other hand, if we start by removing all dominated strategies of Player 1 first, we obtain a solution with payoff $(1, 2)$.

Hence, in the general case of infinite non-zero sum games, the elimination of weakly dominated strategies in infinite games with more than two payoffs may be order dependent and lead to implausible predictions, even if it is not iterated.

Similar phenomena occur in coordination games with infinitely many payoffs. Consider, for instance, the game in Figure 8, where the play ends as soon as one of the players chooses OUT with a payoff corresponding to the number of moves taken in the play. Clearly, any strategy to choose OUT after k many IN-steps is dominated by any strategy that chooses OUT after at least $k+1$ many steps, whenever the current reduction stage contains a strategy of the other player to choose IN for at least k many steps. Therefore, maximal simultaneous elimination would lead both players to choose IN forever and earn payoff 0. Worse than this, a reduction sequence may lead to empty strategy sets. For instance, if we first eliminate all dominated strategies of Player 0, only the one playing IN forever survives. With respect to this single strategy, the strategy of Player 1 to play IN forever is dominated by the one that moves OUT, e.g., at position 2. On the other hand, any strategy to choose OUT at position k (after a sequence of IN) is in turn dominated by the strategy to choose OUT at $k+2$. Thus, none of the strategies available to Player 1 is admissible at this stage of the reduction.

As a larger class of infinte games for which iterated admissibilty may still constitute a meaningful solution concept, the above examples single out coordination games with an arbitrary finite number of payoffs. However, as we shall see, the structure of admissible strategy sets changes substantially

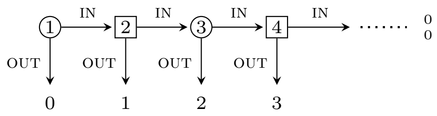

Figure 8: Infinitely many common payoffs

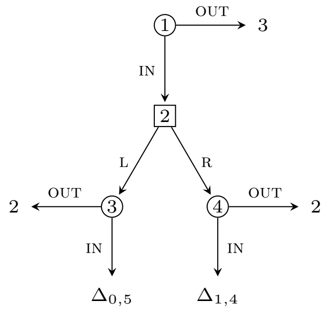

Figure 9: Loss of perfect-information character

when the number of possible payoffs is increased.

Intuitively, for games with two payoffs, the elimination of dominated strategies is backwards inductive, in the sense that any combination of solutions for independent subgames extends to a solution of the entire game, according to the value characterisation in Lemma 3.3. This is no longer the case in games with three or more payoffs. Here, admissibility incorporates a form of forward-inductive reasoning due to which a strategy for a subgame may be dominated or not depending on choices at off-play positions, i.e., positions that are never reached in a play that reaches the subgame. (See [8], for a recent and very thorough account on forwards induction). As a consequence, even though we set out with a game of perfect information, simultaneous maximal elimination of dominated strategies may yield a game restriction which cannot be understood any more as a game with perfect information. Accordingly, partial solutions for independent subgames do not necessarily combine to a solution of the entire game.

To illustrate the effect of forward-inductive reasoning in an infinite coordination game, we adapt an example from [3]. Recall the basic game Δ described in Subsection 3.1 as an example where admissibility cannot ensure coordination, and let us denote by $\Delta_{a,b}$ the modified version where, instead of payoff $+1$ and -1, the players receive payoff a and b, respectively. When considered in isolation, $\Delta_{a,b}$ should of course not be different of $\Delta_{-1,+1}$, for any $a < b$. Let us now look at the game Γ depicted in Figure 9. Here, any strategy of player to choose IN at the root, and OUT both at positions 3 and 4, is dominated by the strategy to choose OUT already at the root. Also, any strategy that is losing on Γ_3 or Γ_4 is (sequentially) dominated. All the other strategies are undominated, because the choices (1IN, 2OUT, 3IN), (1IN, 2IN, 3OUT) and (1IN, 2IN, 3IN), combined with any strategy that is admissible in Γ_3 and Γ_4, lead to mutually incomparable strategies. Hence, it is admissible for Player 0 to play OUT when reaching position 3, only under the counterfactual condition that he would have played IN if the play would have reached position 4. Phrased as a forward-induction argument, this is to say that, upon reaching position 2, Player 1 knows that Player 0 will choose IN either at 3 or at 4, which he could not have derived – by means of admissibility – by considering just the subgames Γ_3 and Γ_4. Thus, in admissible strategies the choices at position 3 and 4 are not any more strategically independent in the sense of [13], although reaching one position would exclude reaching the other one. In fact, the restriction induced by the set of undominated strategies is not equivalent to any game with perfect information.

The loss of perfect information encountered here is in sharp contrast to the case of games with only two payoffs. For infinite games with two – not necessarily common – payoffs it was shown in [4], that strategy sets encountered during simultaneous maximal elimination of dominated strategies satisfy a particular well-formedness criterion: they *allow shifting*, in the phrasing of the paper. Informally, the criterion requires that at any decision point reached during a play, the player i in turn to move at p can deviate from the strategy he used so far and switch to any strategy that is compatible with reaching p. The property that α-admissible strategy sets allow shifting is crucial for most proofs in [4]. As games with more than two payoffs fail to satisfy this property, we cannot rely on our previous results about admissibility in infinite non-zero sum games.

We can neither confirm nor reject the hypothesis that the good properties of admissibility on infinite coordination games with two payoffs are preserved in games with an arbitrary number payoffs. The presence of forward-inductive effects, however, suggests that sharper analytic tools will be needed to settle this question.

References

[1] R. ALUR, T. A. HENZINGER, AND O. KUPFERMAN, *Alternating-time temporal logic*, J. ACM, 49 (2002), pp. 672–713.

[2] K. R. APT, *Uniform proofs of order independence for various strategy elimination procedures*, Contributions to Theoretical Economics, 4 (2004).

[3] P. Battigalli, *On rationalizability in extensive games*, Journal of Economic Theory, 74 (1997), pp. 40–61.

[4] D. Berwanger, *Admissibility in infinite games*, in Proc. of the 24th Annual Symposium on Theoretical Aspects of Computer Science (STACS), vol. 4393 of LNCS, Springer, 2007, pp. 188–199.

[5] C. Bicchieri and O. Schulte, *Common reasoning about admissibility*, Erkenntnis, 45 (1997), pp. 299–325.

[6] D. Gale and F. M. Stewart, *Infinite games with perfect information*, Annals of Mathematics Studies, 28 (1953), pp. 245–266.

[7] I. Gilboa, E. Kalai, and E. Zemel, *On the order of eliminating dominated strategies*, Operation Research Letters, (1988).

[8] S. Govindan and R. Wilson, *On forward induction*, Econometrica, 77 (2009), pp. 1–28.

[9] E. Grädel, W. Thomas, and T. Wilke, eds., *Automata, Logics, and Infinite Games*, vol. 2500 of LNCS, Springer, 2002.

[10] E. Grädel and M. Ummels, *Solution concepts and algorithms for infinite multiplayer games*, in New Perspectives on Games and Interaction, K. Apt and R. van Rooij, eds., vol. 4 of Texts in Logic and Games, Amsterdam University Press, 2008, pp. 151–178.

[11] R. J. Gretlein, *Dominance elimination procedures on finite alternative games*, International Journal of Game Theory, 12 (1983), pp. 107–113.

[12] T. A. Henzinger, *Games in system design and verification.*, in Proc. Theoretical Aspects of Rationality and Knowledge (TARK-2005), 2005.

[13] G. J. Mailath, L. Samuelson, and J. M. Swinkels, *Extensive form reasoning in normal form games*, Econometrica, 61 (1993), pp. 273–302.

[14] L. Marx, *Adaptive learning and iterated weak dominance*, Games and Economic Behavior, 26 (1999), pp. 254–278.

[15] L. M. Marx and J. M. Swinkels, *Order independence for iterated weak dominance*, Games and Economic Behavior, 18 (1997), pp. 219–245.

[16] S. Mohalik and I. Walukiewicz, *Distributed games*, in Proc. 23rd Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS), vol. 2914 of LNCS, Springer, 2003, pp. 338–351.

[17] H. Moulin, *Dominance solvability and Cournot stability*, Mathematical Social Sciences, 7 (1984), pp. 83–102.

[18] L. Samuelson, *Dominated strategies and common knowledge*, Games and Economic Behavior, 4 (1992), pp. 284–313.