Games and Economic Behavior 72 (2011) 510–525

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

An algorithm for proper rationalizability ☆

Andrés Perea

Maastricht University, Department of Quantitative Economics, P.O. Box 616, 6200 MD, Maastricht, The Netherlands

ARTICLE INFO

Article history:
Received 18 December 2009
Available online 27 October 2010

JEL classification:
C72

Keywords:
Epistemic game theory
Proper rationalizability
Algorithms

ABSTRACT

Proper rationalizability (Schuhmacher, 1999; Asheim, 2001) is a concept in epistemic game theory based on the following two conditions: (a) a player should be *cautious*, that is, should not exclude any opponent's strategy from consideration; and (b) a player should *respect the opponents' preferences*, that is, should deem an opponent's strategy s_i infinitely more likely than s_i' if he believes the opponent to prefer s_i to s_i'. A strategy is properly rationalizable if it can optimally be chosen under common belief in the events (a) and (b). In this paper we present an algorithm that for every finite game computes the set of all properly rationalizable strategies. The algorithm is based on the new idea of a *preference restriction*, which is a pair (s_i, A_i) consisting of a strategy s_i, and a subset of strategies A_i, for player i. The interpretation is that player i prefers some strategy in A_i to s_i. The algorithm proceeds by successively adding preference restrictions to the game.

1. Introduction

In a game, it is natural to assume that a player reasons about his opponents before making a decision. Namely, in order to evaluate the possible consequences of a decision, the player must form some belief about his opponents' choices which, in turn, must be based on some belief about his opponents' beliefs about their opponents' choices, and so on. It is the goal of *epistemic game theory* to formally describe such reasoning processes, and to investigate their behavioral implications.

Proper rationalizability (Schuhmacher, 1999; Asheim, 2001) is a concept within epistemic game theory that is based upon the following two assumptions:

- a player should be *cautious*, that is, a player should not exclude any opponent's strategy from consideration;
- a player should *respect the opponents' preferences,* that is, if the player believes that an opponent prefers strategy s_i to strategy s_i', then the player should deem s_i much more likely (in fact, *infinitely* more likely) than s_i'.

Any strategy that can be chosen optimally under common belief in these two events is called *properly rationalizable*.

In order to define proper rationalizability formally we can no longer model the players' beliefs by standard probability distributions. Suppose, for instance, that player 1 believes that player 2 prefers strategy a to strategy b. If player 1's belief about 2's choice would be modeled by a single probability distribution then player 1 should assign probability 0 to b, since he must respect 2's preferences. This, however, would contradict the assumption that he is cautious.

A possible way to define proper rationalizability is by means of *sequences of probability distributions*, as Schuhmacher (1999) does, or by using *lexicographic probability systems*, as Asheim (2001) does. Both frameworks can model a state of

☆ I would like to thank Geir Asheim, Amanda Friedenberg, Jean-Jacques Herings and Oliver Schulte for interesting discussions and valuable comments. I am also indebted to some anonymous referees who have provided useful suggestions that improved the paper.

E-mail address: a.perea@maastrichtuniversity.nl.
URL: http://www.personeel.unimaas.nl/a.perea/.

0899-8256/$ – see front matter
doi:10.1016/j.geb.2010.10.008

mind in which you deem some opponent's strategy s_i infinitely more likely than some other strategy s_i', without completely discarding the latter choice.

The practical disadvantage of these richer frameworks is that, in many examples, it makes the computation of properly rationalizable strategies rather difficult. This is probably also the reason that proper rationalizability, despite its strong intuitive appeal, has not received as much attention as many other concepts in game theory. It would therefore be very useful to have an algorithm helping us to compute these properly rationalizable strategies. Schuhmacher (1999) presents a procedure, called *iteratively proper trembling*, that for any given $\varepsilon > 0$ yields the set of ε-properly rationalizable strategies. By letting ε tend to zero, we finally would obtain the set of properly rationalizable strategies. So, in a sense, Schuhmacher's procedure only *indirectly* leads to the set of properly rationalizable strategies, as we first have to apply the procedure for a sequence of small ε's, and then let ε go to zero.

Schulte (2003) provides another algorithm designed for proper rationalizability, called *iterated backward inference*. This procedure does not exactly yield the set of properly rationalizable strategies, as its output may contain strategies that are not properly rationalizable. The output, however, always *includes* the set of properly rationalizable strategies.

In this paper we present an algorithm, called *iterated addition of preference restrictions*, that *directly* delivers the set of all properly rationalizable strategies in every finite game. The algorithm is based on the new notion of a *preference restriction*. Formally, a preference restriction for player i is a pair (s_i, A_i), where s_i is a strategy and A_i a subset of strategies for player i. The interpretation is that player i prefers some strategy in A_i to s_i, without specifying which one (unless A_i contains only one strategy, of course). A *lexicographic belief* for player i about his opponents' strategies is a finite sequence $\lambda_i = (\lambda_i^1, \ldots, \lambda_i^K)$ of probability distributions on S_{-i}, the set of opponents' strategy combinations, such that every strategy combination s_{-i} in S_{-i} receives positive probability under some probability distribution λ_i^k in this sequence. For every $k \in \{1, \ldots, K\}$, we call λ_i^k the level k belief. The lexicographic belief λ_i deems some strategy combination s_{-i} *infinitely more likely* than some other strategy combination s'_{-i} if there is some level k such that s_{-i} receives positive probability under the level k belief λ_i^k, whereas s'_{-i} receives probability zero under the first k levels. We say that λ_i *respects* a preference restriction (s_j, A_j) for opponent j if it deems some strategy in A_j infinitely more likely than s_j. This thus mimics the condition in proper rationalizability that i must respect j's preferences. The lexicographic belief λ_i is said to *assume* a subset $D_{-i} \subseteq S_{-i}$ of strategy combinations if it deems every element in D_{-i} infinitely more likely than every element outside D_{-i} (cf. Brandenburger et al., 2008).

The algorithm we present proceeds by inductively adding preference restrictions, until no further preference restrictions can be produced. At round 1, we start with the empty set of preference restrictions for all players. In every subsequent round, we add a preference restriction (s_i, A_i) for player i if every lexicographic belief on S_{-i} that respects all current preference restrictions for i's opponents, assumes some subset $D_{-i} \subseteq S_{-i}$ on which s_i is weakly dominated by some randomized strategy on A_i. We continue this process until no further preference restriction can be added. Among the final set of preference restrictions for player i, we look for those strategies s_i that are not part of any preference restriction (s_i, A_i). We show that these strategies are exactly the properly rationalizable strategies for player i.

So, at every round the algorithm produces, for each player, a set of preference restrictions. As the set of preference restrictions can only grow at every round, and there are only finitely many possible preference restrictions, the algorithm must stop after finitely many rounds.

Not only can this algorithm be used to *compute* the properly rationalizable strategies in a game, it also represents a natural *inductive reasoning procedure* for the players, that eventually leads them to properly rationalizable choices. The central object in this reasoning process is that of a preference restriction. If we add a preference restriction (s_i, A_i) for player i, then epistemically this means that i's opponents believe that i prefers some strategy in A_i to s_i. Moreover, if i's opponents respect i's preferences, as we assume in proper rationalizability, then i's opponents will also deem some strategy in A_i infinitely more likely than s_i. Thus, by adding preference restrictions at every round, we further and further restrict the possible lexicographic beliefs that players can plausibly hold about their opponents' choices. In a sense, what the algorithm shows is that, in order to reason your way toward properly rationalizable strategies, it is sufficient to keep track of the players' preference restrictions. At every round, by considering the current preference restrictions, we can possibly derive new preference restrictions, thus further restricting the players' possible lexicographic beliefs, until this reasoning process cannot produce any new preference restrictions. This is where the reasoning procedure ends, and by looking at the final preference restrictions we can find all the properly rationalizable strategies in the game.

In the algorithm we present, the objects of output are different than in Schuhmacher's procedure. There, the procedure delivers at every round, and for every player i, a set of full support probability distributions on player i's strategies, where this set becomes smaller with every round. As there are infinitely many possible sets of full support probability distributions, Schuhmacher's procedure can produce infinitely many possible outputs in every round. This is a major difference with the algorithm we propose, where at every round there is only a finite number of possible outputs, namely the preference restrictions at that round.

Note also that the algorithm in this paper is fundamentally different from most other inductive concepts in epistemic game theory, which usually proceed by successively eliminating strategies from the game. Think, for instance, of iterated elimination of strictly (weakly) dominated strategies, and the Dekel–Fudenberg procedure (Dekel and Fudenberg, 1990) consisting of one round of elimination of weakly dominated strategies, followed by iterated elimination of strictly dominated strategies. So, why did we not base the algorithm on elimination of strategies as well? The reason is that iterated elimina-

	d	e	f
a	0, 0	0, 1	3, 1
b	1, 1	1, 0	1, 1
c	2, 1	2, 1	2, 0

Fig. 1. Why elimination of strategies does not work.

tion of strategies cannot work for proper rationalizability. In Section 2 we provide an example that shows this. Hence, an algorithm for proper rationalizability must necessarily be of a different nature than the ones we are used to.

The outline of the paper is as follows. In Section 2 we show, by means of an example, why successive elimination of strategies does not work for proper rationalizability. In Section 3 we give a formal definition of proper rationalizability, by making use of lexicographic probability systems. In Section 4 we present the algorithm, illustrate it by means of an example, and state our main theorem showing that the algorithm produces exactly the set of properly rationalizable strategies. In Section 5 we discuss some important properties of the algorithm: We show how the algorithm can be viewed as a natural inductive reasoning procedure, and explain why the order in which we add preference restrictions does not matter for the eventual output. Section 6, finally, contains all the proofs.

2. Why elimination of strategies does not work

Most algorithms in the epistemic game theory literature proceed by successively eliminating *strategies* from the game. Think, for instance, of iterated elimination of strictly (weakly) dominated strategies, and the Dekel–Fudenberg procedure (Dekel and Fudenberg, 1990) consisting of one round of elimination of weakly dominated strategies, followed by iterated elimination of strictly dominated strategies. As announced in the introduction, the algorithm we propose for proper rationalizability is of a different nature since it is based on successively adding *preference restrictions* rather than eliminating *strategies*. A natural question is why we do not stick to the process of eliminating strategies here. In this section we show why elimination of strategies does not work for proper rationalizability.

Let us first be precise about the class of strategy elimination procedures we consider. All the elimination procedures mentioned above have in common that at each round, only weakly dominated strategies in the reduced game (but not necessarily all) are eliminated. Now, say that a strategy elimination procedure is *regular* if at every round, it eliminates a (possibly empty) subset of the set of weakly dominated strategies in the reduced game.

We will show, by means of an example, that a regular strategy elimination procedure cannot work for proper rationalizability. Consider the game in Fig. 1, where player 1 is the row player and player 2 the column player. Let us first see what proper rationalizability does for this example. Since player 1 prefers c to b, player 2 must deem c infinitely more likely than b. But then, player 2 will prefer e to f, and hence player 1 must deem e infinitely more likely than f. This, in turn, implies that player 1 prefers b to a, and therefore player 2 must deem b infinitely more likely than a. So, overall, player 2 must deem c infinitely more likely than b, and b infinitely more likely than a. As a consequence, player 2 must choose d, and player 1 must choose c. Hence, proper rationalizability uniquely selects strategy c for player 1, and strategy d for player 2.

We now show that a regular strategy elimination procedure can never eliminate strategy e for player 2. Note that there is only one weakly dominated strategy in the full game, namely b. So, at the first round we either eliminate nothing, or we eliminate strategy b for player 1. However, if we eliminate b, then strategy e can never become weakly dominated in any smaller game, so we would never be able to eliminate strategy e after eliminating b. Hence, we see that by applying a regular strategy elimination procedure, we will never eliminate strategy e from the game, despite the fact that e is not properly rationalizable. Such a procedure can therefore not work for proper rationalizability.

The key problem here is that, according to proper rationalizability, player 2 must deem b infinitely more likely than a (see our argument above). However, at the same time, a regular strategy elimination procedure can only eliminate strategy b at the beginning, which amounts to requiring that player 2 must deem the remaining strategies, a and c, infinitely more likely than b. This, obviously, produces a conflict.

3. Definition of proper rationalizability

The concept of *proper rationalizability* has first been defined by Schuhmacher (1999). More precisely, Schuhmacher introduces for every $\varepsilon > 0$ the ε-proper trembling condition as an analogue to Myerson's (1978) condition underlying proper equilibrium. For a given ε, the concept of ε-proper rationalizability is formalized by imposing common belief in the ε-proper trembling condition. Proper rationalizability is obtained, finally, by letting ε approach 0. Although Schuhmacher provides, for every $\varepsilon > 0$, an epistemic model for ε-proper rationalizability, he does not give a direct epistemic foundation for the limiting concept of proper rationalizability. Later, Asheim (2001) has provided an epistemic foundation for the limiting concept of proper rationalizability in two-player games, making use of *lexicographic beliefs*. In this section, we use Asheim's model and extend it to more than two players.

3.1. Lexicographic probability systems

Lexicographic probability systems have been formally introduced by Blume et al. (1991a, 1991b) as a possible way to represent a decision maker's belief about the state of the world. The essential feature is that it allows the decision maker to deem one state much more likely (in fact, infinitely more likely) than some other state, without completely ignoring the latter state when making a decision.

More formally, let X be some finite set of states. By $\Delta(X)$ we denote the set of all probability distributions on X. A *lexicographic probability system* (LPS) on X is a finite sequence of probability distributions

$$\lambda = (\lambda^1, \lambda^2, \ldots, \lambda^K),$$

with $\lambda^k \in \Delta(X)$ for all $k \in \{1, \ldots, K\}$. We refer to λ^1 as the decision maker's *level* 1 *belief*, to λ^2 as his *level* 2 *belief*, and so on. The interpretation is that the decision maker attaches much more importance to his level 1 belief than to his level 2 belief, attaches much more importance to his level 2 belief than to his level 3 belief, and so on, without completely discarding any of these beliefs. For every state $x \in X$, let $rk(x, \lambda)$ be the first level k for which $\lambda^k(x) > 0$. If $\lambda^k(x) = 0$ for every $k \in \{1, \ldots, K\}$, set $rk(x, \lambda) = \infty$. We call $rk(x, \lambda)$ the *rank* of state x within the LPS λ. We say that the LPS λ deems state x *infinitely more likely* than some other state y if x has a lower rank that y.

3.2. Epistemic model

Consider a finite static game $\Gamma = (S_i, u_i)_{i\in I}$ where I is the finite set of players, the finite set S_i denotes the set of strategies for player i, and $u_i : \prod_{j\in I} S_j \to \mathbb{R}$ denotes player i's utility function. We assume that player i does not only have a belief about his opponents' strategy choices, but also about the possible beliefs that his opponents could have about the other players' strategy choices, and about the possible beliefs that the opponents could have about the possible beliefs that their opponents could have about the other players' strategy choices, and so on. That is, player i holds a full *belief hierarchy* about the opponents' choices and the opponents' beliefs. If we assume, moreover, that each of the beliefs in this hierarchy can be represented by an LPS, this leads to the following epistemic model.

Definition 3.1 *(Epistemic model).* A finite epistemic model for the game Γ is a tuple $(T_i, \lambda_i)_{i\in I}$ where, for all players i, T_i is a finite set of types, and λ_i is a function that assigns to every type $t_i \in T_i$ some LPS $\lambda_i(t_i)$ on the set $S_{-i} \times T_{-i}$ of opponents' strategy–type combinations.

Here, $S_{-i} := \prod_{j\neq i} S_j$ denotes the set of opponents' strategy combinations, and $T_{-i} := \prod_{j\neq i} T_j$ the set of opponents' type combinations. The interpretation is that $\lambda_i(t_i)$ represents the belief that type t_i has about his opponents' choices and beliefs. For instance, the marginal of $\lambda_i(t_i)$ on S_j represents the belief that t_i has about opponent j's choice. Since every opponent's type t_j holds a belief about the other players' choices, we can derive from $\lambda_i(t_i)$ as well the belief that type t_i has about the belief that player j has about his opponents' choices, and so on. In fact, from $\lambda_i(t_i)$ we can derive the full belief hierarchy that player i has about his opponents' choices and beliefs.

The reader may wonder why we restrict attention to epistemic models with *finitely* many types for every player. The reason is that this is sufficient for the purpose of this paper. In principle, we could allow for infinitely many types for every player, and define proper rationalizability for such infinite epistemic models. But it can be shown that every properly rationalizable strategy in a finite game can be supported by a properly rationalizable type within an epistemic model with *finitely* many types only. So, we do not "overlook" any properly rationalizable strategies by concentrating on finite type spaces only. As working with finite sets of types makes things easier, we have decided to solely concentrate on finite epistemic models in this paper.

Note that within an epistemic model, the lexicographic belief $\lambda_i(t_i) = (\lambda_i^1, \ldots, \lambda_i^K)$ of a type t_i is, mathematically speaking, an LPS on the set of states $S_{-i} \times T_{-i}$. For every opponents' strategy–type combination $(s_{-i}, t_{-i}) \in S_{-i} \times T_{-i}$, we can thus define the *rank* $rk((s_{-i}, t_{-i}), \lambda_i(t_i))$ of (s_{-i}, t_{-i}) within $\lambda_i(t_i)$, being the lowest level k such that $\lambda_i^k(s_{-i}, t_{-i}) > 0$. Remember that, by convention, $rk((s_{-i}, t_{-i}), \lambda_i(t_i)) = \infty$ whenever (s_{-i}, t_{-i}) does not receive positive probability anywhere in $\lambda_i(t_i)$. We say that type t_i deems the strategy–type combination (s_{-i}, t_{-i}) *infinitely more likely* than some other combination (s'_{-i}, t'_{-i}) if the rank of (s_{-i}, t_{-i}) is lower than the rank of (s'_{-i}, t'_{-i}).

Similarly, we can define for every *event* $E \subseteq S_{-i} \times T_{-i}$ of opponents' strategy–type combinations the associated rank by

$$rk(E, \lambda_i(t_i)) = \min\{r((s_{-i}, t_{-i}), \lambda_i(t_i)) \mid (s_{-i}, t_{-i}) \in E\}.$$

Hence, the rank of E is the lowest level k such that λ_i^k assigns positive probability to some element in E. This definition then allows us to define the rank of an *individual* opponent's strategy–type pair (s_j, t_j), simply by taking the rank of the event

$$\{s_j\} \times \prod_{k\neq i,j} S_k \times \{t_j\} \times \prod_{k\neq i,j} T_k.$$

So, we first take the marginal of the LPS $\lambda_i(t_i)$ on $S_j \times T_j$, and then take the rank of (s_j, t_j) inside this marginal LPS. In a similar fashion, we can also define the rank of an individual opponent's type t_j, and of an individual opponent's strategy s_j. As such, we can formally state expressions like "$\lambda_i(t_i)$ deems (s_j, t_j) infinitely more likely than (s'_j, t'_j) for opponent j" or "$\lambda_i(t_i)$ deems s_j infinitely more likely than s'_j for opponent j", which means that the rank of the former is smaller than the rank of the latter.

We say that type t_i *deems possible* some event $E \subseteq S_{-i} \times T_{-i}$ if there is some level k with $\lambda_i^k(E) > 0$. That is, E is deemed possible if and only if $rk(E, \lambda_i(t_i)) \neq \infty$. Since we have defined the rank also for individual strategy–type pairs (s_j, t_j) and for individual types t_j, we can also formally define the event that type t_i deems possible a strategy–type pair (s_j, t_j) for opponent j, and that t_i deems possible an opponent's type t_j. It simply means that the associated rank is not ∞.

3.3. Cautious types

Intuitively, *caution* means that the player should not fully exclude any opponent's choice from consideration. The formal definition is a little bit more subtle, however – it states that a type t_i should not exclude any strategy choice *for any opponent's type t_j he considers possible*. Hence, for every belief hierarchy that t_i deems possible for his opponent j, and for every strategy s_j that j can possibly choose, type t_i should deem possible the event that his opponent holds this belief hierarchy *and* chooses s_j.

Definition 3.2 *(Cautious type).* Consider an epistemic model with sets of types T_i for every player i. Type $t_i \in T_i$ is cautious if, for every opponent j, every type $t_j \in T_j$ he considers possible, and every strategy choice $s_j \in S_j$, type t_i deems possible the strategy–type pair (s_j, t_j).

3.4. Respecting the opponents' preferences

The key condition in Asheim's model for proper rationalizability is that a type should *respect his opponents' preferences.* In words it means that, whenever type t_i believes that his opponent j prefers some strategy s_j to some other strategy s'_j, then he should deem s_j infinitely more likely than s'_j. We must first define what it means, within our epistemic model, that a type prefers some strategy to another strategy.

Consider a type t_i with an LPS $\lambda_i(t_i) = (\lambda_i^1, \dots, \lambda_i^K)$ on $S_{-i} \times T_{-i}$. Then, for every level $k \in \{1, \dots, K\}$ and every strategy s_i, we can define the *level k expected utility*

$$u_i\big(s_i, \lambda_i^k\big) := \sum_{(s_{-i},t_{-i}) \in S_{-i} \times T_{-i}} \lambda_i^k(s_{-i}, t_{-i})\, u_i(s_i, s_{-i}).$$

This is the expected utility that would result by choosing s_i under the belief λ_i^k.

Definition 3.3 *(A type's preference relation over strategies).* Let $t_i \in T_i$ be a type with LPS $\lambda_i(t_i) = (\lambda_i^1, \dots, \lambda_i^K)$ on $S_{-i} \times T_{-i}$. Type t_i prefers strategy s_i to some other strategy s'_i if there is some level $k \in \{1, \dots, K\}$ such that $u_i(s_i, \lambda_i^k) > u_i(s'_i, \lambda_i^k)$ and $u_i(s_i, \lambda_i^l) = u_i(s'_i, \lambda_i^l)$ for all $l < k$.

For later purposes, we say that type t_i *weakly prefers s_i to s'_i* if t_i does not prefer s'_i to s_i.

Definition 3.4 *(Respecting the opponents' preferences).* Let $t_i \in T_i$ be a cautious type. Type t_i respects the opponent's preferences if, for every opponent j, every type $t_j \in T_j$ deemed possible by t_i, and every two strategies s_j, s'_j such that t_j prefers s_j to s'_j, type t_i deems the pair (s_j, t_j) infinitely more likely than the pair (s'_j, t_j).

3.5. Proper rationalizability

We say that a type t_i is *properly rationalizable* if t_i is cautious and respects the opponents' preferences, believes that all opponents are cautious and respect their opponents' preferences, believes that all opponents believe that their opponents are cautious and respect their opponents' preferences, and so on. In other words, t_i is cautious and respects the opponents' preferences, and expresses common belief in the event that players are cautious and respect the opponents' preferences.

Definition 3.5 *(Common belief in "caution and respect of the opponents' preferences").* A type t_i expresses common belief in the event that players are cautious and respect the opponents' preferences if t_i only deems possible opponents' types that are cautious and respect their opponents' preferences, only deems possible opponents' types that only deem possible opponents' types that are cautious and respect their opponents' preferences, and so on.

By additionally assuming that t_i itself is cautious and respects the opponents' preferences, we obtain the definition of a properly rationalizable type.

Definition 3.6 *(Properly rationalizable type).* A type t_i is properly rationalizable if it is cautious and respects the opponents' preferences, and moreover expresses common belief in the event that players are cautious and respect the opponents' preferences.

Finally, we say that a strategy s_i is properly rationalizable for player i if it is optimal for some properly rationalizable type. Formally, a strategy s_i is called *optimal* for type t_i if t_i weakly prefers s_i to any other strategy.

Definition 3.7 *(Properly rationalizable strategy).* A strategy s_i for player i is properly rationalizable if there is some finite epistemic model $(T_i, \lambda_i)_{i \in I}$ and some properly rationalizable type $t_i \in T_i$ such that s_i is optimal for t_i.

As we already mentioned before, the concept of a properly rationalizable strategy would not change if we would allow for infinite epistemic models here.

4. Algorithm

In this section we will present an algorithm that always delivers all properly rationalizable strategies. Before doing so, we first provide some intuitive arguments that eventually will lead to the algorithm. We will then state the algorithm formally, and illustrate it by means of an example. Finally, we state our main result, namely that the algorithm yields precisely the set of properly rationalizable strategies in every game. The proof for this result can be found in Section 6.

4.1. Road to the algorithm

In Section 2 we have seen that elimination of (subsets of) weakly dominated strategies cannot work for proper rationalizability. So, what kind of procedure *could* work here? We start our informal investigation with the following well known fact:

Step 1. Suppose that strategy s_i is weakly dominated on S_{-i} by some randomized strategy $\mu_i \in \Delta(A_i)$, where A_i is a subset of strategies. Then, if player i is cautious, he will prefer some strategy in A_i to s_i. We say that (s_i, A_i) is a *preference restriction* for player i.

Here, $\Delta(A_i)$ denotes the set of probability distributions on A_i. The reason for this fact is simple: If s_i is weakly dominated by μ_i, then under every cautious lexicographic belief, s_i will be worse than μ_i, and hence there must be some $a_i \in A_i$ which is better than s_i under such a cautious lexicographic belief. So, (s_i, A_i) will be a preference restriction for player i.

Suppose now that player i believes his opponents are cautious, and that he respects his opponents' preferences. If some opponent's strategy s_j is weakly dominated on S_{-j} by some randomized strategy $\mu_j \in \Delta(A_j)$, then we know by Step 1 that player j will prefer some strategy in A_j to s_j in case he is cautious. As player i indeed believes he is cautious, and respects j's preferences, player i must deem some strategy in A_j infinitely more likely than s_j. We say that player i's lexicographic belief *respects* the preference restriction (s_j, A_j). This leads to the following observation:

Step 2. Suppose player i believes his opponents are cautious, and respects his opponents' preferences. Then, i's lexicographic belief must respect every opponent's preference restriction (s_j, A_j) generated in Step 1.

Say that a lexicographic belief for player i *assumes* a set $D_{-i} \subseteq S_{-i}$ of opponents' strategy combinations if it deems all strategy combinations inside D_{-i} infinitely more likely than all strategy combinations outside D_{-i}. Suppose now that i's lexicographic belief is cautious, and assumes some set D_{-i} of opponents' strategy combinations. Assume, moreover, that his strategy s_i is weakly dominated on D_{-i} by a randomized strategy $\mu_i \in \Delta(A_i)$. Then, i must prefer some strategy in A_i to s_i. The argument is basically the same as for Step 1, if we would "reduce" the game to opponents' strategy combinations in D_{-i}. We thus obtain the following step:

Step 3. Suppose that every lexicographic belief for player i respecting all preference restrictions from Step 1, assumes some $D_{-i} \subseteq S_{-i}$ on which s_i is weakly dominated by some $\mu_i \in \Delta(A_i)$. Suppose, moreover, that player i is cautious, believes his opponents are cautious, and respects the opponents' preferences. Then, i must prefer some strategy in A_i to s_i. We say that (s_i, A_i) is a *new preference restriction* for player i.

Of course, we can iterate this argument if we assume that player i is cautious, respects the opponents' preferences, and expresses common belief in the event that players are cautious and respect the opponents' preferences. That is, if we assume that player i's type is properly rationalizable. The inductive step would then look as follows:

Inductive step. Suppose that every lexicographic belief for i that respects all preference restrictions generated so far, assumes some $D_{-i} \subseteq S_{-i}$ on which s_i is weakly dominated by some $\mu_i \in \Delta(A_i)$. Then, if i is of a properly rationalizable type, he must prefer some strategy in A_i to s_i. So, (s_i, A_i) would be a new preference restriction for player i.

This would thus generate an inductive procedure in which at every step (possibly) some new preference restrictions would be added for the players. Since there are only finitely many possible preference restrictions for the players, this procedure must end after finitely many steps. Now, consider some player i, and his set of preference restrictions generated by the procedure above. If player i is of some properly rationalizable type, we know from our arguments above that he will never choose a strategy s_i if it is part of some preference restriction (s_i, A_i). In that case, namely, he would always prefer some strategy in A_i to s_i, so s_i could not be optimal.

So, the procedure above rules out strategies that are certainly not properly rationalizable. But what about the converse? So, what about strategies that are not ruled out by the procedure above? The main theorem in this paper, Theorem 4.6, will show that the "surviving" strategies are all properly rationalizable! Hence, the procedure above will always select *exactly* those strategies that are properly rationalizable – not more and not less.

4.2. Description of the algorithm

Before we state the algorithm, we first formally define the new concepts we described above, such as preference restrictions, what it means for a lexicographic belief to respect a preference restriction, and so on.

Definition 4.1 *(Preference restriction).* A preference restriction for player i is a pair (s_i, A_i) where s_i is a strategy, and A_i a nonempty subset of strategies.

The interpretation is that player i prefers at least one strategy from A_i to s_i. Now, consider a lexicographic belief λ_i on S_{-i}, which is simply an LPS on S_{-i}. From here on, we will always assume that such a lexicographic belief λ_i has full support on S_{-i}, that is, every strategy combination in S_{-i} receives positive probability in some level of λ_i.

Definition 4.2 *(Respecting a preference restriction).* A lexicographic belief λ_i on S_{-i} respects a preference restriction (s_j, A_j) for player j if λ_i deems some strategy in A_j infinitely more likely than s_j.

This, in a sense, mimics the requirement that player i must respect j's preferences.

Definition 4.3 *(Assuming a set of opponents' strategy combinations).* Consider a subset $D_{-i} \subseteq S_{-i}$ of opponents' strategy combinations, and a lexicographic belief λ_i on S_{-i}. The lexicographic belief λ_i assumes the set D_{-i} if λ_i deems all strategy combinations inside D_{-i} infinitely more likely than all strategy combinations outside D_{-i}.

This notion is based upon the idea of "assuming an event" in Brandenburger et al. (2008). Note that a lexicographic belief $\lambda_i = (\lambda_i^1, \ldots, \lambda_i^K)$ on S_{-i} assumes a subset $D_{-i} \subseteq S_{-i}$, if and only if, there is some level $k \in \{1, \ldots, K\}$ such that $\bigcup_{l \leqslant k} \text{supp}(\lambda_i^l) = D_{-i}$. Here, $\text{supp}(\lambda_i^l)$ denotes the support of the probability distribution λ_i^l.

A *randomized strategy* for player i is a probability distribution $\mu_i \in \Delta(S_i)$ on player i's strategies. For a subset $A_i \subseteq S_i$, we denote by $\Delta(A_i)$ the set of randomized strategies that assign positive probability only to strategies in A_i. For some opponents' strategy combination $s_{-i} \in S_{-i}$, let

$$u_i(\mu_i, s_{-i}) := \sum_{s_i \in S_i} \mu_i(s_i)\, u_i(s_i, s_{-i})$$

denote i's expected utility from the randomized strategy μ_i and the opponents' strategy combination s_{-i}.

Definition 4.4 *(Weakly dominated strategy).* Let $D_{-i} \subseteq S_{-i}$ be a subset of the opponents' strategy combinations. Strategy s_i is said to be weakly dominated by randomized strategy μ_i on D_{-i} if $u_i(\mu_i, s_{-i}) \geqslant u_i(s_i, s_{-i})$ for all $s_{-i} \in D_{-i}$, with strict inequality for at least some $s_{-i} \in D_{-i}$.

We are now ready to present the algorithm. The idea is to start with the empty set of preference restrictions for all players, and at every round to add new preference restrictions, if possible. For that reason, the algorithm is called "iterated addition of preference restrictions".

Algorithm 4.5 *(Iterated addition of preference restrictions).* In round 1, begin for all players i with the empty set of preference restrictions.

At every further round $n \geqslant 2$, restrict for every player i to those lexicographic beliefs on S_{-i} that respect all opponents' preference restrictions generated so far. Add a new preference restriction (s_i, A_i) for player i if every such lexicographic belief assumes some set $D_{-i} \subseteq S_{-i}$ on which s_i is weakly dominated by some $\mu_i \in \Delta(A_i)$.

Since the number of preference restrictions is finite, this algorithm must end after a finite number of rounds. We say that strategy s_i *survives* the algorithm of iterated addition of preference restrictions if s_i is not part of any preference

	e	f	g	h
a	3, 3	0, 0	1, 1	0, 0
b	0, 0	3, 3	1, 1	0, 0
c	1, 0	1, 0	1, 1	5, 0
d	4, 0	4, 0	4, 0	4, 0

Fig. 2. Illustration of the algorithm.

restriction (s_i, A_i) generated by the algorithm. Namely, if s_i were to be part of a preference restriction (s_i, A_i) produced by the algorithm, then player i would prefer at least one strategy in A_i to s_i, and hence s_i could not be optimal.

4.3. Illustration of the algorithm

We will now illustrate the algorithm by means of an example. Consider the game from Fig. 2.

Round 1. We start with the empty set of preference restrictions for both players.

Round 2. Clearly, every lexicographic belief for player 2 assumes the set $\{a, b, c, d\}$. Since h is weakly dominated by e, f and g on $\{a, b, c, d\}$, we add the preference restrictions

$$\big(h, \{e\}\big), \big(h, \{f\}\big) \text{ and } \big(h, \{g\}\big)$$

for player 2.

Round 3. Every lexicographic belief for player 1 that respects the preference restrictions $(h, \{e\}), (h, \{f\})$ and $(h, \{g\})$ must deem e, f and g infinitely more likely than h, and hence must assume the set $\{e, f, g\}$. On $\{e, f, g\}$, strategies a, b and c are weakly dominated by d, and c is weakly dominated by the randomized strategy $\frac{1}{2}a + \frac{1}{2}b$. So, we add the preference restrictions

$$\big(a, \{d\}\big), \big(b, \{d\}\big), \big(c, \{d\}\big) \text{ and } \big(c, \{a, b\}\big)$$

for player 1.

Round 4. Every lexicographic belief for player 2 that respects the preference restriction $(c, \{a, b\})$ must deem a or b infinitely more likely than c. So, every such belief must assume some set $D_1 \subseteq S_1$ that contains a or b, but not c. On every such set D_1, strategy g is weakly dominated by the randomized strategy $\frac{1}{2}e + \frac{1}{2}f$, and hence we add the preference restriction

$$\big(g, \{e, f\}\big)$$

for player 2.

After this round no new preference restrictions can be generated, apart from those that are "logically implied" by the ones above. By this, we mean the following: If we take a preference restriction (s_i, A_i), then it logically implies all the preference restrictions $(s_i, \hat{A}_i)$ with $A_i \subseteq \hat{A}_i$.

So, the algorithm generates the preference restrictions

$$\big(a, \{d\}\big), \big(b, \{d\}\big), \big(c, \{d\}\big) \text{ and } \big(c, \{a, b\}\big)$$

for player 1, and the preference restrictions

$$\big(h, \{e\}\big), \big(h, \{f\}\big), \big(h, \{g\}\big) \text{ and } \big(g, \{e, f\}\big)$$

for player 2, plus those that are logically implied by these. For player 1, the only strategy s_1 that is not part of a preference restriction (s_1, A_1) is strategy d. For player 2, the only strategies s_2 that are not part of a preference restriction (s_2, A_2) are e and f. Hence, the strategies that survive iterated addition of preference restrictions are d for player 1, and e and f for player 2. We show that d, e and f are exactly the properly rationalizable strategies in the game.

Consider the epistemic model as given in Table 1. This table should be read as follows: We consider two types for player 1, $\{t_1, \hat{t}_1\}$, and two types for player 2, $\{t_2, \hat{t}_2\}$. Type t_1 only deems possible opponent's type t_2, and deems the strategy–type pair (e, t_2) infinitely more likely than the strategy–type pair (g, t_2), which he deems infinitely more likely than (f, t_2), which, in turn, he deems infinitely more likely than (h, t_2). Similarly for the other types in the model.

It can easily be verified that every type in this model is cautious and respects the opponent's preferences. Therefore, every type in this model expresses common belief in the event that both players are cautious and respect the opponent's preferences. This implies that every type in this model is properly rationalizable. As strategy d is optimal for t_1 and $\hat{t}_1$,

Table 1
An epistemic model for the game in Fig. 2.

Types	$T_1 = \{t_1, \hat{t}_1\}$, $T_2 = \{t_2, \hat{t}_2\}$
Beliefs for player 1	$\lambda_1(t_1) = ((e, t_2), (g, t_2), (f, t_2), (h, t_2))$ $\lambda_1(\hat{t}_1) = ((f, \hat{t}_2), (g, \hat{t}_2), (e, \hat{t}_2), (h, \hat{t}_2))$
Beliefs for player 2	$\lambda_2(t_2) = ((d, t_1), (a, t_1), (c, t_1), (b, t_1))$ $\lambda_2(\hat{t}_2) = ((d, \hat{t}_1), (b, \hat{t}_1), (c, \hat{t}_1), (a, \hat{t}_1))$

strategy e is optimal for t_2, and strategy f is optimal for $\hat{t}_2$, we conclude that d, e and f are indeed properly rationalizable strategies in this game.[1]

The reader may verify that there are no other properly rationalizable strategies in this game. As such, d, e and f are the only properly rationalizable strategies in the game. So, in this example, the algorithm yields exactly the properly rationalizable strategies for all players. Our main theorem in this paper states that this is always the case!

4.4. Main theorem

Our main theorem states that the algorithm of iterated addition of preference restrictions yields *exactly* the set of properly rationalizable strategies for every player.

Theorem 4.6 *(Algorithm yields precisely the set of properly rationalizable strategies). Consider a finite static game. Then, a strategy s_i is properly rationalizable, if and only if, s_i survives the algorithm of iterated addition of preference restrictions.*

The proof for this result can be found in Section 6. The easier direction is to show that every properly rationalizable strategy survives iterated addition of preference restriction. So, a properly rationalizable strategy s_i can never be part of a preference restriction (s_i, A_i) generated by the algorithm. The proof for this direction is basically a formalization of the intuitive arguments laid out at the beginning of this section. The more difficult direction is to prove that every strategy s_i that is not part of any such preference restriction (s_i, A_i) is properly rationalizable. Hence, we must construct an epistemic model in which each of these strategies s_i is supported by some properly rationalizable type. This construction is rather delicate.

From the theorem, we can easily derive the following observation: If in a given game *no* strategy is weakly dominated, then *all* strategies for the players are properly rationalizable. Namely, the algorithm we present will only generate preference restrictions at the first round if there is at least some strategy that is weakly dominated within the full game. Otherwise, the algorithm will not generate any preference restriction at all, and hence all strategies would survive the algorithm.

4.5. A finite formulation of the algorithm

The algorithm of iterated addition of preference restrictions as we have formulated it, proceeds by adding preference restrictions and deleting lexicographic beliefs at every round. More precisely, we start with the empty set of preference restrictions and the full set of lexicographic beliefs. At the first round we see whether we can add some preference restrictions. If so, then this would reduce the set of lexicographic beliefs, which at the next round could add some further preference restrictions, and so on.

What is somewhat undesirable from a computational point of view is that there are infinitely many possible lexicographic beliefs in the game. This would suggest that at every round in the algorithm we must scan through infinitely many lexicographic beliefs. This, however, is not necessary. What matters for the algorithm is not so much the precise probabilities in the lexicographic belief, but the induced "likelihood ordering" on opponents' strategy combinations. More precisely, let $\lambda_i = (\lambda_i^1, \ldots, \lambda_i^K)$ be a lexicographic belief on S_{-i}. Remember our convention that λ_i has full support on S_{-i}, that is, every $s_{-i} \in S_{-i}$ receives positive probability in some level λ_i^k. Let $L_i = (L_i^1, \ldots, L_i^M)$ be the ordered sequence of disjoint subsets $L_i^m \subseteq S_{-i}$ such that (a) λ_i deems every $s_{-i} \in L_i^m$ infinitely more likely than every $s'_{-i} \in L_i^{m+1}$, for every $m \in \{1, \ldots, M-1\}$, (b) for every m and every $s_{-i}, s'_{-i} \in L_i^m$, the LPS λ_i does not deem s_{-i} infinitely more likely than s'_{-i}, nor vice versa, and (c) the union of the sets in L_i is S_{-i}. We call L_i the *likelihood ordering* induced by λ_i. Formally, we have the following definition.

Definition 4.7 *(Likelihood ordering).* A likelihood ordering for player i on the opponents' strategy combinations is an ordered sequence $L_i = (L_i^1, \ldots, L_i^M)$ where $L_i^1, \ldots, L_i^M$ are pairwise disjoint subsets of S_{-i} whose union is equal to S_{-i}.

[1] In fact, (d, e) and (d, f) are proper *equilibria* (Myerson, 1978) in this game. Proper equilibria correspond to pairs (t_1, t_2) of properly rationalizable types such that t_1 only deems possible the opponent's type t_2, and t_2 only deems possible the opponent's type t_1. (This follows from Blume et al., 1991b, Proposition 5.) Note that the pairs (t_1, t_2) and $(\hat{t}_1, \hat{t}_2)$ in Table 1 have this property, and they correspond to the proper equilibria (d, e) and (d, f), respectively. However, in general there are properly rationalizable strategies that cannot be supported by type pairs (t_1, t_2) that have the property above. This is because not every properly rationalizable strategy is part of a proper equilibrium.

So, the interpretation is that L_i deems all strategy combinations in L_i^1 infinitely more likely than all strategy combinations in L_i^2, deems all strategy combinations in L_i^2 infinitely more likely than all strategy combinations in L_i^3, and so on. It is clear that there are only finitely many likelihood orderings in the game, since there are only finitely many strategies for every player.

We can now easily extend the definitions of "respecting a preference restriction" and "assuming a set of opponents' strategy combinations" to likelihood orderings. Say that a likelihood ordering $L_i = (L_i^1, \ldots, L_i^M)$ *respects a preference restriction* (s_j, A_j) if L_i deems some strategy in A_j infinitely more likely than s_j. Also, the likelihood ordering L_i is said to *assume the set* D_{-i} *of opponents' strategy combinations* if L_i deems all strategy combinations inside D_{-i}, infinitely more likely than all strategy combinations outside D_{-i}.

The algorithm of iterated addition of preference restrictions can thus alternatively be stated as follows:

Algorithm 4.8 *(Finite version).* In round 1, begin for all players i with the empty set of preference restrictions.

At every further round $n \geqslant 2$, restrict for every player i to those likelihood orderings on S_{-i} that respect all opponents' preference restrictions generated so far. Add a new preference restriction (s_i, A_i) for player i if every such likelihood ordering assumes some set $D_{-i} \subseteq S_{-i}$ on which s_i is weakly dominated by some $\mu_i \in \Delta(A_i)$.

The advantage of this formulation is that at every round, we only have to scan through finitely many objects, as there are only finitely many preference restrictions and likelihood orderings in the game. Obviously, this algorithm generates precisely the same set of preference restrictions as the original procedure. As such, the properly rationalizable strategies are precisely those strategies that survive this alternative algorithm.

5. Discussion

In this section we will discuss some important properties of the algorithm.

5.1. Algorithm as an inductive reasoning procedure

The algorithm is not merely a tool to compute the properly rationalizable strategies in a game, but can also be interpreted as an inductive reasoning process that can be used by a player who reasons in the spirit of proper rationalizability. Consider namely a fixed player in the game, say player i. In round 2, the algorithm would add for every opponent j a preference restriction (s_j, A_j) if s_j would be weakly dominated on S_{-j} by a mixture on A_j. In that case, player i would store the preference restriction (s_j, A_j) in his mind, meaning that he believes that player j prefers some strategy in A_j to s_j. If i respects j's preferences, then he should consequently deem some strategy in A_j infinitely more likely than s_j. That is, the preference restrictions that player i would store in his mind at round 2 would restrict the possible lexicographic beliefs he could hold about his opponents' choices. Moreover, if player i believes that his opponents reason similarly, then player i can actually deduce the possible lexicographic beliefs that his opponents may hold at this round.

In the next round of his reasoning procedure, player i would then ask for every opponent j: Given his restricted set of beliefs, would player j always assume some set $D_{-j} \subseteq S_{-j}$ on which some strategy s_j would always be weakly dominated by a mixture on A_j? If yes, then player i will store (s_j, A_j) as a new preference restriction in his mind. By doing so, player i would then further restrict the possible lexicographic beliefs he could hold about his opponents. Player i could continue this inductive reasoning procedure until no new preference restriction could be added, and hence his possible lexicographic beliefs could not be restricted any further.

So we see that the algorithm may serve very well as an intuitive reasoning procedure for players, that will eventually lead them to the properly rationalizable strategies in the game. What is crucial in this reasoning procedure is that a player only needs to keep track of preference restrictions, which substantially simplifies matters compared to the original definition of proper rationalizability. In that light, our main theorem thus says that in order to find the properly rationalizable strategies in a game, it is sufficient for a player to think in terms of preference restrictions, and to reason in accordance with the algorithm.

In the epistemic game theory literature, there are other algorithms that can nicely be interpreted as intuitive reasoning procedures. Take, for instance, the epistemic concept of *common belief in rationality* (Tan and Werlang, 1988) and the associated algorithm of *iterated elimination of strictly dominated strategies.* Here, the algorithm can be seen as an epistemic reasoning procedure in which a player successively deletes opponents' strategies from his mind, since they can no longer be optimal. At every round, this would then restrict the player's possible beliefs as he must assign probability zero to these strategies. These additional restrictions on the players' beliefs could then induce further strategies that can be deleted, and so on. So, in that procedure the players' possible (non-lexicographic) beliefs are restricted further and further by deleting strategies, whereas in our procedure the (lexicographic) beliefs are restricted further and further by adding new preference restrictions.

A similar story can be told for the epistemic concept of *iterated assumption of rationality within a complete type structure* (Brandenburger et al., 2008) and the associated algorithm of *iterated elimination of weakly dominated strategies.* Here, the algorithm reflects an epistemic reasoning procedure in which a player with lexicographic beliefs iteratedly deletes weakly dominated strategies from his mind. At every round of this procedure, the player will then deem all surviving strategies infinitely more likely than all deleted strategies, thus restricting the possible lexicographic beliefs he can hold (see Stahl, 1995, who proposes exactly this type of reasoning). So also in this procedure, the player's possible beliefs are restricted in every round by deleting strategies.

Other algorithms that can be interpreted as epistemic reasoning procedures are, for instance, the *Dekel–Fudenberg procedure* (Dekel and Fudenberg, 1990) for static games, and *extensive form rationalizability* (Pearce, 1984; Battigalli, 1997) and *backwards induction* (Zermelo, 1913) for dynamic games.

5.2. Order independence

For the algorithm, it can be shown that the order and speed in which we add preference restrictions does not matter for the eventual result. That is, it does not matter whether in every round we add *all* preference restrictions that can possibly be generated, or only *some* of these.

To see this, let us compare two procedures, Procedure 1 and Procedure 2, where in the first we always add *all* possible preference restrictions at every round, and in the second we only add *some* of the possible preference restrictions every time. Then, first of all, Procedure 1 will at every round generate at least as many preference restrictions as Procedure 2. Namely, at round 2 Procedure 1 generates as least as many preference restrictions, by definition. Therefore, at round 3 Procedure 1 restricts to a smaller set of lexicographic beliefs than Procedure 2. But then, under Procedure 1 it will be "easier" to generate new preference restrictions at round 3 than under Procedure 2. Hence, at round 3 Procedure 1 will, again, generate at least as many preference restrictions as Procedure 2, and so on. So, eventually, Procedure 1 will generate as least as many preference restrictions as Procedure 2. The key argument here was that a larger set of preference restrictions will lead to a smaller set of possible lexicographic beliefs, and a smaller set of possible lexicographic beliefs will in turn lead to a larger set of induced preference restrictions. So, the algorithm is *monotone* in this sense.

On the other hand, it can also be shown that every preference restriction generated by Procedure 1 will also eventually be generated by Procedure 2. Suppose, namely, that Procedure 1 would generate some preference restriction that would not be generated at all by Procedure 2. Then, let k be the first round at which Procedure 1 would generate a preference restriction, say (s_i, A_i), not generated by Procedure 2 at all. By construction of the algorithm, every lexicographic belief for player i that respects all preference restrictions generated by Procedure 1 *before* round k, must assume some set D_{-i} on which s_i is weakly dominated by some $\mu_i \in \Delta(A_i)$. By our assumption, all these preference restrictions generated by Procedure 1 before round k are also eventually generated by Procedure 2, let us say before round $m \geqslant k$. But then, every lexicographic belief for player i that respects all preference restrictions generated by Procedure 2 before round m, assumes a set D_{-i} on which s_i is weakly dominated by some $\mu_i \in \Delta(A_i)$. Hence, Procedure 2 must add the preference restriction (s_i, A_i) sooner or later, which is a contradiction since we assumed that Procedure 2 does not generate preference restriction (s_i, A_i) at all. We thus conclude that every preference restriction added by Procedure 1 is also finally added by Procedure 2. As such, Procedures 1 and 2 eventually generate exactly the same set of preference restrictions. So, indeed, the order and speed in which we add preference restrictions is irrelevant to the algorithm.

6. Proofs

In this section we prove the main theorem (Theorem 4.6), stating that the algorithm of iterated addition of preference restrictions selects exactly the set of properly rationalizable strategies in the game. We start by laying out three preparatory results that will be useful for proving the main theorem.

6.1. Preparatory results

For our first preparatory result, we recall the definition of a *likelihood ordering induced by an LPS* as we gave it in Section 4.5. Consider an LPS $\lambda_i = (\lambda_i^1, \dots, \lambda_i^K)$ on S_{-i}. Remember our convention that λ_i has full support on S_{-i}, that is, every $s_{-i} \in S_{-i}$ receives positive probability in some level λ_i^k. Let $L_i = (L_i^1, \dots, L_i^M)$ be the ordered sequence of disjoint subsets $L_i^m \subseteq S_{-i}$ such that (a) λ_i deems every $s_{-i} \in L_i^m$ infinitely more likely than every $s'_{-i} \in L_i^{m+1}$, for every $m \in \{1, \dots, M-1\}$, (b) for every m and every $s_{-i}, s'_{-i} \in L_i^m$, the LPS λ_i does not deem s_{-i} infinitely more likely than s'_{-i}, nor vice versa, and (c) the union of the sets in L_i is S_{-i}. We call L_i the *likelihood ordering* induced by λ_i. Our first result characterizes, for a given strategy s_i and set $A_i \subseteq S_i$, those likelihood orderings on S_{-i} that admit an LPS under which s_i is weakly preferred to all strategies in A_i.

Lemma 6.1. *Let λ_i be an LPS on S_{-i}, let s_i be a strategy and $A_i \subseteq S_i$ a subset of strategies.*

(a) *If under the LPS λ_i, strategy s_i is weakly preferred to all strategies in A_i, then λ_i does not assume any $D_{-i} \subseteq S_{-i}$ on which s_i is weakly dominated by a mixture on A_i.*
(b) *If λ_i does not assume any $D_{-i} \subseteq S_{-i}$ on which s_i is weakly dominated by a mixture on A_i, then there is some LPS σ_i, inducing the same likelihood ordering as λ_i, under which s_i is weakly preferred to all strategies in A_i.*

This result is actually a generalization of Lemma 4 in Pearce (1984) which shows that a strategy s_i is not weakly dominated if and only if it is optimal for a full support probability distribution on S_{-i}. Take, namely, an LPS λ_i on S_{-i} with one level only, and choose $A_i = S_i$. By our definition of an LPS on S_{-i}, the belief λ_i has full support on S_{-i}. So, the only subset D_{-i} that is assumed by λ_i is S_{-i}. Moreover, every LPS σ_i inducing the same likelihood ordering as λ_i must be a single level LPS with full support on S_{-i}, so must be a single full support probability distribution on S_{-i}. But then, for such choices of λ_i and A_i, our lemma says that (a) every s_i that is optimal under λ_i must not be weakly dominated on S_{-i}, and (b) every s_i that is not weakly dominated on S_{-i} is optimal for some full support probability distribution σ_i on S_{-i}. This is precisely Pearce's result.

Proof of Lemma 6.1. (a) Suppose that under the LPS $\lambda_i = (\lambda_i^1, \ldots, \lambda_i^K)$, strategy s_i is weakly preferred to all strategies in A_i. Assume, contrary to what we want to prove, that λ_i assumes some $D_{-i} \subseteq S_{-i}$ on which s_i is weakly dominated by some $\mu_i \in \Delta(A_i)$. As λ_i assumes D_{-i}, we know from Section 4.2 that there must be some $k \in \{1, \ldots, K\}$ with $\bigcup_{l \leqslant k} \text{supp}(\lambda_i^l) = D_{-i}$. Since μ_i weakly dominates s_i on D_{-i}, we have that $u_i(s_i, \lambda_i^l) \leqslant u_i(\mu_i, \lambda_i^l)$ for all $l \leqslant k$, with strict inequality for at least some $l \leqslant k$. Here, $u_i(s_i, \lambda_i^l)$ denotes the expected utility of choosing s_i under the belief λ_i^l, and $u_i(\mu_i, \lambda_i^l)$ denotes the expected utility of $\mu_i \in \Delta(A_i)$ under λ_i^l. This means that for every $l \leqslant k$, either (1) $u_i(s_i, \lambda_i^l) = u_i(a_i, \lambda_i^l)$ for every $a_i \in A_i$, or (2) there is some $a_i \in A_i$ with $u_i(s_i, \lambda_i^l) < u_i(a_i, \lambda_i^l)$. Moreover, case (2) must apply for at least one $l \leqslant k$. This implies that there is some $a_i \in A_i$ that is preferred to s_i under λ_i. However, this is a contradiction to our assumption above that s_i is weakly preferred to all strategies in A_i under λ_i. Hence, we conclude that λ_i cannot assume a subset D_{-i} on which s_i is weakly dominated by some mixture on A_i.

(b) Suppose that $\lambda_i = (\lambda_i^1, \ldots, \lambda_i^K)$ does not assume any $D_{-i} \subseteq S_{-i}$ on which s_i is weakly dominated by a mixture on A_i. Let $L_i = (L_i^1, \ldots, L_i^M)$ be the induced likelihood ordering. Then, the subsets D_{-i} that are assumed by λ_i are exactly the sets $L_i^1 \cup \cdots \cup L_i^m$ for $m \in \{1, \ldots, M\}$. Hence, for every $m \in \{1, \ldots, M\}$ strategy s_i is not weakly dominated by any $\mu_i \in \Delta(A_i)$ on $L_i^1 \cup \cdots \cup L_i^m$. By Lemma 4 in Pearce (1984), there is for every $m \in \{1, \ldots, M\}$ some probability distribution $\sigma_i^m \in \Delta(L_i^1 \cup \cdots \cup L_i^m)$, with full support on $L_i^1 \cup \cdots \cup L_i^m$, such that under σ_i^m strategy s_i is weakly preferred to all strategies in A_i. But then, $\sigma_i := (\sigma_i^1, \ldots, \sigma_i^M)$ is an LPS inducing the same likelihood ordering as λ_i, namely $L_i = (L_i^1, \ldots, L_i^M)$. Moreover, under σ_i strategy s_i is weakly preferred to all strategies in A_i. This completes the proof. □

Our second preparatory result shows that a properly rationalizable type always respects all the preference restrictions generated by the algorithm. To state and prove this result formally, we need the following notation. For every round n in the algorithm of *iterated addition of preference restrictions*, let R_i^n be the set of preference restrictions generated for player i at round n. Similarly, let R_{-i}^n be the set of preference restrictions generated for i's opponents at round n. That is, (s_i, A_i) is a preference restriction in R_i^n if and only if every lexicographic belief on S_{-i} respecting all preference restrictions in R_{-i}^{n-1}, assumes some D_{-i} on which s_i is weakly dominated by some mixture on A_i. By R_i^∞ and R_{-i}^∞ we denote the sets of all preference restrictions for player i, and i's opponents, that have been generated when the algorithm stops.

Lemma 6.2. *Let t_i be a properly rationalizable type. Then, t_i's lexicographic belief on S_{-i} respects every preference restriction in R_{-i}^∞.*

Proof. We show that for all n, all players i, and every properly rationalizable type t_i, the lexicographic belief that t_i holds on S_{-i} respects every preference restriction in R_{-i}^n. We prove this by induction on n.

For $n = 1$ the statement is trivial since R_{-i}^1 is the empty set of preference restrictions, and hence every lexicographic belief on S_{-i} respects all preference restrictions in R_{-i}^1.

Now, let $n \geqslant 2$ and suppose that, for all players i and every properly rationalizable type t_i, the belief of t_i on S_{-i} respects every preference restriction in R_{-i}^{n-1}. Take a properly rationalizable type t_i. We prove that t_i's belief respects every preference restriction in R_{-i}^n.

As type t_i is properly rationalizable, t_i only considers possible opponents' types t_j that are properly rationalizable. By the induction assumption, it follows that t_i only considers possible opponents' types t_j that respect every preference restriction in R_{-j}^{n-1}.

Take an opponent j, and a preference restriction $(s_j, A_j) \in R_j^n$. Then, by construction of the algorithm, every lexicographic belief for player j that respects all preference restrictions in R_{-j}^{n-1} must assume some $D_{-j} \subseteq S_{-j}$ on which s_j is weakly dominated by some mixture on A_j. By part (a) in Lemma 6.1, it follows that under every lexicographic belief for player j that respects all preference restrictions in R_{-j}^{n-1}, player j prefers some strategy in A_j to s_j. Since we have seen that t_i only considers possible types t_j that respects all preference restrictions in R_{-j}^{n-1}, we may conclude that type t_i only considers possible types t_j that prefer some strategy in A_j to s_j.

Since t_i is properly rationalizable, it respects the opponents' preferences, and hence t_i must deem some strategy in A_j infinitely more likely than s_j. Summarizing, we have seen that for every preference restriction $(s_j, A_j) \in R_j^n$, type t_i deems some strategy in A_j infinitely more likely than s_j. This, however, means that t_i respects all preference restrictions in R_{-i}^n, which was to be shown. By induction, the proof is complete. □

The third lemma describes an important property of the sets of preference restrictions that are *not* generated by the algorithm. This result will be crucial for proving that every strategy that survives the algorithm, that is, is not part of any preference restriction produced by the algorithm, is properly rationalizable. It will be the basis, namely, for constructing our properly rationalizable types. For this lemma, we need the following notation: For a given LPS λ_i on S_{-i}, and an opponent's strategy s_j, we denote by $A_j^-(s_j, \lambda_i)$ the set of strategies for player j that are not deemed infinitely more likely than s_j by λ_i. Hence, $A_j^-(s_j, \lambda_i)$ contains those strategies that receive equal, or higher, rank than s_j under the LPS λ_i.

Lemma 6.3 *(Property of preference restrictions not generated by the algorithm). For every player i, let R_i^{not} be the set of preference restrictions not generated by the algorithm. Then, for every $(s_i, A_i) \in R_i^{not}$ there is an LPS λ_i on S_{-i} such that*

(1) *under λ_i, strategy s_i is weakly preferred to all strategies in A_i, and*
(2) *for every opponent's strategy s_j, the pair $(s_j, A_j^-(s_j, \lambda_i))$ is in R_j^{not}.*

Proof. Let $(s_i, A_i) \in R_i^{not}$. So, (s_i, A_i) is not generated by the algorithm, that is, $(s_i, A_i) \notin R_i^\infty$. Then, by construction of the algorithm, there is some lexicographic belief λ_i' on S_{-i}, respecting all preference restrictions in R_{-i}^∞, that does not assume any D_{-i} on which s_i is weakly dominated by some mixture on A_i. By Lemma 6.1, for every such λ_i' there is a lexicographic belief λ_i, inducing the same likelihood ordering as λ_i', under which s_i is weakly preferred to all strategies in A_i. But then, since λ_i and λ_i' induce the same likelihood ordering, also λ_i respects all preference restrictions in R_{-i}^∞. Hence, for (s_i, A_i) there is some lexicographic belief λ_i on S_{-i}, respecting all preference restrictions in R_{-i}^∞, under which s_i is weakly preferred to all strategies in A_i. This proves (1).

Now, take an opponent's strategy s_j. By definition, λ_i deems no strategy in $A_j^-(s_j, \lambda_i)$ infinitely more likely than s_j. As λ_i respects all preference restrictions in R_j^∞, it must thus be the case that $(s_j, A_j^-(s_j, \lambda_i))$ is not in R_j^∞, and hence $(s_j, A_j^-(s_j, \lambda_i))$ is in R_j^{not}. This proves (2). □

Our last preparatory result provides a method of "blowing up" an LPS on S_{-i} without changing the induced preference relation on S_i.

Lemma 6.4 *(Blow up lemma). Let $\lambda_i = (\lambda_i^1, \ldots, \lambda_i^K)$ and $\sigma_i = (\sigma_i^1, \ldots, \sigma_i^L)$ be two LPS's on S_{-i} such that for every $k \in \{1, \ldots, K\}$ there is some $l(k) \in \{1, \ldots, L\}$ with the property that (1) $\sigma_i^{l(k)} = \lambda_i^k$, and (2) $\sigma_i^m \in \{\lambda_i^1, \ldots, \lambda_i^{k-1}\}$ for every $m < l(k)$. So, σ_i can be seen as a "blown up" version of λ_i.*

Then, λ_i and σ_i induce the same preference relation on S_i.

Proof. Take some strategies $s_i, s_i' \in S_i$. We show that s_i is preferred to s_i' under λ_i if and only if s_i is preferred to s_i' under σ_i.

(a) Suppose that s_i is preferred to s_i' under λ_i. Then, there is some $k \in \{1, \ldots, K\}$ such that $u_i(s_i, \lambda_i^k) > u_i(s_i', \lambda_i^k)$ and $u_i(s_i, \lambda_i^m) = u_i(s_i', \lambda_i^m)$ for all $m < k$. But then, $u_i(s_i, \sigma_i^{l(k)}) > u_i(s_i', \sigma_i^{l(k)})$ and $u_i(s_i, \sigma_i^l) = u_i(s_i', \sigma_i^l)$ for all $l < l(k)$. So, s_i is preferred to s_i' under σ_i.

(b) Suppose that s_i is preferred to s_i' under σ_i. Then, there is some $l \in \{1, \ldots, L\}$ such that $u_i(s_i, \sigma_i^l) > u_i(s_i', \sigma_i^l)$ and $u_i(s_i, \sigma_i^m) = u_i(s_i', \sigma_i^m)$ for all $m < l$. Consequently, there is no copy of σ_i^l that appears before level l in σ_i, and hence $l = l(k)$ for some $k \in \{1, \ldots, K\}$. This implies that $u_i(s_i, \lambda_i^k) > u_i(s_i', \lambda_i^k)$ and $u_i(s_i, \lambda_i^m) = u_i(s_i', \lambda_i^m)$ for all $m < k$, and hence s_i is preferred to s_i' under λ_i. This completes the proof. □

6.2. Proof of the main theorem

In this subsection we prove our main theorem (Theorem 4.6), which states that a strategy is properly rationalizable if and only if it survives the procedure of iterated addition of preference restrictions. We thus must prove two directions: (a) Every properly rationalizable strategy survives the procedure of iterated addition of preference restrictions, and (b) Every strategy that survives this procedure is properly rationalizable. As we will see, (b) is the more difficult direction to prove.

Proof of (a): Every properly rationalizable strategy survives the algorithm. As before, let R_i^∞ be the final set of preference restrictions for player i, and R_{-i}^∞ the final set of preference restrictions for i's opponents, generated by the algorithm. Take a properly rationalizable strategy s_i for player i. We must show that s_i is not part of any preference restriction (s_i, A_i) in R_i^∞.

Since s_i is properly rationalizable, there is a properly rationalizable type t_i for which s_i is optimal. That is, for type t_i strategy s_i is weakly preferred to all strategies in S_i. Hence, by part (a) in Lemma 6.1, type t_i's lexicographic belief on S_{-i} does not assume any D_{-i} on which s_i is weakly dominated by some mixture on S_i. Moreover, as t_i is properly rationalizable, it follows from Lemma 6.2 that t_i respects all preference restrictions in R_{-i}^∞. Summarizing, we thus see that t_i's lexicographic belief on S_{-i} respects all preference restrictions in R_{-i}^∞, but does not assume any D_{-i} on which s_i is weakly dominated by a mixture on S_i. But then, by construction of the algorithm, s_i cannot be part of any preference restriction (s_i, A_i) in R_i^∞. So, s_i survives the algorithm.

Proof of (b): Every strategy that survives the algorithm is properly rationalizable. As before, let R_i^{not} be the set of preference restrictions for player i that are *not* generated by the algorithm. Then, by Lemma 6.3 we know that for every (s_i, A_i) in R_i^{not} there is an LPS $\lambda_i(s_i, A_i)$ on S_{-i} such that

- under $\lambda_i(s_i, A_i)$, strategy s_i is weakly preferred to all strategies in A_i, and
- for every opponent's strategy s_j, the pair $(s_j, A_j^-(s_j, \lambda_i(s_i, A_i)))$ is in R_j^{not}.

Remember that $A_j^-(s_j, \lambda_i(s_i, A_i))$ contains those strategies that receive equal, or higher, rank than s_j under the LPS $\lambda_i(s_i, A_i)$.

The idea will now be to construct, for every $(s_i, A_i) \in R_i^{not}$, some properly rationalizable type $t_i(s_i, A_i)$. We define, for every player i, the set of types

$$T_i = \{t_i(s_i, A_i) \mid (s_i, A_i) \in R_i^{not}\}.$$

Our task will be to assign to every type $t_i(s_i, A_i)$ an LPS $\sigma_i(s_i, A_i)$ on $S_{-i} \times T_{-i}$ such that:

- every $\sigma_i(s_i, A_i)$ induces the same preference relation on S_i as $\lambda_i(s_i, A_i)$ does,
- every $\sigma_i(s_i, A_i)$ is cautious, and
- every $\sigma_i(s_i, A_i)$ respects the opponents' preferences.

Suppose we would have completed this task. Then, first of all, every type in T_i would be properly rationalizable, since it would be cautious and respect the opponents' preferences, and consider possible only opponents' types in T_{-i} which are all cautious and respect the opponents' preferences, and so on.

Next, consider a strategy s_i that survives the algorithm, that is, which is not part of a preference restriction (s_i, A_i) in R_i^∞. Then, (s_i, S_i) is in R_i^{not}, and hence $t_i(s_i, S_i) \in T_i$. By construction, under the LPS $\lambda_i(s_i, S_i)$ strategy s_i is weakly preferred to all strategies in S_i. Since type $t_i(s_i, S_i)$ holds the LPS $\sigma_i(s_i, S_i)$, and $\sigma_i(s_i, S_i)$ induces the same preference relation on S_i as $\lambda_i(s_i, S_i)$, strategy s_i is optimal for type $t_i(s_i, S_i)$. As $t_i(s_i, S_i)$ is properly rationalizable, we conclude that strategy s_i is properly rationalizable. Hence, every strategy s_i that survives the algorithm would be properly rationalizable. This would complete the proof of the main theorem.

So, if we can carry out our task above, the proof would be complete. We construct the LPS's $\sigma_i(s_i, A_i)$ in two steps. In Step 1, we assign to every type $t_i(s_i, A_i)$ an LPS $\rho_i(s_i, A_i)$ on $S_{-i} \times T_{-i}$ such that:

- every $\rho_i(s_i, A_i)$ induces the same preference relation on S_i as $\lambda_i(s_i, A_i)$ does, but
- $\rho_i(s_i, A_i)$ is not yet cautious.

We shall refer to the belief levels in $\rho_i(s_i, A_i)$ as the "main levels". In Step 2, we make a blown up version $\sigma_i(s_i, A_i)$ of the LPS $\rho_i(s_i, A_i)$, by adding "blow up levels" in the sense of Lemma 6.4, such that:

- every $\sigma_i(s_i, A_i)$ induces the same preference relation on S_i as $\rho_i(s_i, A_i)$ does,
- every $\sigma_i(s_i, A_i)$ is cautious, and
- every $\sigma_i(s_i, A_i)$ respects the opponents' preferences.

Step 1 *(Construction of main levels).* Fix a pair (s_i, A_i) in R_i^{not}. Consider the associated LPS $\lambda_i(s_i, A_i) = (\lambda_i^1, \ldots, \lambda_i^K)$ on S_{-i}. Recall that under $\lambda_i(s_i, A_i)$ strategy s_i is weakly preferred to all strategies in A_i, and that for every opponent's strategy s_j, the pair $(s_j, A_j^-(s_j, \lambda_i(s_i, A_i)))$ is in R_j^{not}. So, for every opponent's strategy s_j, we have that $t_j(s_j, A_j^-(s_j, \lambda_i(s_i, A_i)))$ is a type in T_j. To reduce notation, let us from now on write $A_j^-(s_j)$ instead of $A_j^-(s_j, \lambda_i(s_i, A_i))$, since we will fix the LPS $\lambda_i(s_i, A_i)$.

The LPS $\rho_i(s_i, A_i) = (\rho_i^1, \ldots, \rho_i^K)$ on $S_{-i} \times T_{-i}$ is then defined as follows: For every $k \in \{1, \ldots, K\}$,

$$\rho_i^k\big((s_j, t_j)_{j \neq i}\big) := \begin{cases} \lambda_i^k((s_j)_{j \neq i}), & \text{if } t_j = t_j(s_j, A_j^-(s_j)) \text{ for every } j \neq i \\ 0, & \text{otherwise} \end{cases}$$

for every $(s_j, t_j)_{j \neq i} \in S_{-i} \times T_{-i}$. Hence, ρ_i^k induces probability distribution λ_i^k on S_{-i} for every k. Consequently, $\rho_i(s_i, A_i)$ induces the same preference relation on S_i as $\lambda_i(s_i, A_i)$ does. Moreover, $\rho_i(s_i, A_i)$ only deems possible strategy–type pairs $(s_j, t_j(s_j, A_j^-(s_j)))$ where $s_j \in S_j$.

Step 2 *(Construction of blow up levels).* Take an LPS $\rho_i(s_i, A_i)$ constructed above. Note that $\rho_i(s_i, A_i)$ is not cautious: For every opponents' type $t_j(s_j, A_j^-(s_j))$ it considers possible, there is only one strategy it considers possible, namely s_j. So, in order to extend $\rho_i(s_i, A_i)$ to a cautious LPS, we need to add extra belief levels that cover all pairs $(s_j', t_j(s_j, A_j^-(s_j)))$ with $s_j' \neq s_j$.

For every opponent j, and every pair (s_j', s_j) with $s_j' \neq s_j$, we define a "blow up" level $\tau_i(s_j', s_j) \in \Delta(S_{-i} \times T_{-i})$ as follows: Let k be the first level such that ρ_i^k assigns positive probability to s_j'. Then, $\tau_i(s_j', s_j)$ is a copy of ρ_i^k, except for the fact that

$\tau_i(s_j', s_j)$ shifts the probability that ρ_i^k assigned to the pair $(s_j', t_j(s_j', A_j^-(s_j')))$ completely toward the pair $(s_j', t_j(s_j, A_j(s_j)))$. In particular, $\tau_i(s_j', s_j)$ induces the same probability distribution on S_{-i} as ρ_i^k.

Without loss of generality, let us fix an opponent j and a strategy $s_j \in S_j$. Let l be the first level such that ρ_i^l assigns positive probability to s_j. Suppose that the LPS $\lambda_j(s_j, A_j^-(s_j))$ induces the ordering $(s_j^1, \ldots, s_j^M)$ on S_j, meaning that under $\lambda_j(s_j, A_j^-(s_j))$ strategy s_j^1 is weakly preferred to s_j^2, that s_j^2 is weakly preferred to s_j^3, and so on. Suppose further that $s_j = s_j^m$, and that all strategies $s_j^1, \ldots, s_j^{m-1}$ are strictly preferred to s_j.

We insert blow up levels $\tau_i(s_j^1, s_j), \ldots, \tau_i(s_j^{m-1}, s_j)$ between main levels ρ_i^{l-1} and ρ_i^l in this particular order. So, $\tau_i(s_j^1, s_j)$ comes before $\tau_i(s_j^2, s_j)$, and so on. We then insert blow up levels $\tau_i(s_j^{m+1}, s_j), \ldots, \tau_i(s_j^M, s_j)$ after the last main level ρ_i^K in this particular order.

If we do so for every opponent j and every strategy $s_j \in S_j$, we obtain a cautious LPS $\sigma_i(s_i, A_i)$ with main levels ρ_i^k and blow up levels $\tau_i(s_j', s_j)$ in between. This completes the construction of the LPS's $\sigma_i(s_i, A_i)$ for every (s_i, A_i) in R_i^{not}.

Step 3 *(Every LPS $\sigma_i(s_i, A_i)$ induces the same preference relation on S_i as $\rho_i(s_i, A_i)$).* We now prove that every LPS $\sigma_i(s_i, A_i)$ so constructed induces the same preference relation on S_i as $\rho_i(s_i, A_i)$. By construction, the main levels in $\sigma_i(s_i, A_i)$ coincide exactly with the levels $\rho_i^1, \ldots, \rho_i^K$ in $\rho_i(s_i, A_i)$. Consider a blow up level $\tau_i(s_j', s_j)$ that comes before main level ρ_i^k. We show that $\tau_i(s_j', s_j)$ induces the same probability distribution on S_{-i} as some ρ_i^m with $m < k$.

By our construction above, s_j must receive positive probability in some ρ_i^l with $l \leqslant k$, and under $\lambda_j(s_j, A_j^-(s_j))$ strategy s_j' must be preferred to s_j. Since, by definition of $\lambda_j(s_j, A_j^-(s_j))$, strategy s_j is weakly preferred to every strategy in $A_j^-(s_j)$ under $\lambda_j(s_j, A_j^-(s_j))$, it must be that $s_j' \notin A_j^-(s_j)$. By definition, $A_j^-(s_j)$ contains all those strategies that are not deemed infinitely more likely than s_j by $\lambda_i(s_i, A_i)$, and hence s_j' must be deemed infinitely more likely than s_j by $\lambda_i(s_i, A_i)$. By construction of $\rho_i(s_i, A_i)$, this implies that $\rho_i(s_i, A_i)$ deems s_j' infinitely more likely than s_j. Since s_j receives positive probability in some ρ_i^l with $l \leqslant k$, strategy s_j' receives positive probability for the first time in some ρ_i^m with $m < k$. But then, by construction of $\tau_i(s_j', s_j)$, the blow up level $\tau_i(s_j', s_j)$ is a copy of ρ_i^m, except for the fact that $\tau_i(s_j', s_j)$ shifts the probability that ρ_i^m assigned to the pair $(s_j', t_j(s_j', A_j^-(s_j')))$ completely toward the pair $(s_j', t_j(s_j, A_j^-(s_j)))$. In particular, $\tau_i(s_j', s_j)$ induces the same probability distribution on S_{-i} as ρ_i^m. So, we have shown that every blow up level $\tau_i(s_j', s_j)$ that comes before main level ρ_i^k induces the same probability distribution on S_{-i} as some main level ρ_i^m with $m < k$.

Now, let $\hat{\sigma}_i(s_i, A_i)$ be the marginal of $\sigma_i(s_i, A_i)$ on S_{-i}, and let $\hat{\rho}_i(s_i, A_i) = (\hat{\rho}_i^1, \ldots, \hat{\rho}_i^K)$ be the marginal of $\rho_i(s_i, A_i)$ on S_{-i}. Let $\hat{\tau}_i(s_j', s_j)$ be the marginal of the blow up level $\tau_i(s_j', s_j)$ on S_{-i}. By our insight above, we may conclude that every blow up level $\hat{\tau}_i(s_j', s_j)$ that comes before main level $\hat{\rho}_i^k$ in $\hat{\sigma}_i(s_i, A_i)$ is a copy of some $\hat{\rho}_i^m$ with $m < k$. This means, however, that $\hat{\sigma}_i(s_i, A_i)$ is a blown up version of $\hat{\rho}_i(s_i, A_i)$ in the sense of Lemma 6.4, and hence, by the same lemma, $\hat{\sigma}_i(s_i, A_i)$ induces the same preference relation on S_i as $\hat{\rho}_i(s_i, A_i)$. Consequently, $\sigma_i(s_i, A_i)$ induces the same preference relation on S_i as $\rho_i(s_i, A_i)$, which was to show.

Step 4 *(Every LPS $\sigma_i(s_i, A_i)$ respects the opponents' preferences).* We will now show that every $\sigma_i(s_i, A_i)$ respects the opponent's preferences. Suppose that $\sigma_i(s_i, A_i)$ deems possible some opponents' type $t_j(s_j, A_j)$, and that $t_j(s_j, A_j)$ prefers s_j' to s_j''. We show that $\sigma_i(s_i, A_i)$ deems $(s_j', t_j(s_j, A_j))$ infinitely more likely than $(s_j'', t_j(s_j, A_j))$.

Since $t_j(s_j, A_j)$ is deemed possible by $\sigma_i(s_i, A_i)$, it must be the case that $t_j(s_j, A_j) = t_j(s_j, A_j^-(s_j))$. By construction, type $t_j(s_j, A_j^-(s_j))$ holds LPS $\sigma_j(s_j, A_j^-(s_j))$ which, we have seen, induces the same preference relation on S_j as $\lambda_j(s_j, A_j^-(s_j))$. Since, by assumption, the type $t_j(s_j, A_j^-(s_j))$ prefers s_j' to s_j'', it follows that s_j' is preferred to s_j'' under $\lambda_j(s_j, A_j^-(s_j))$. But then, the construction of the blow up levels in $\sigma_i(s_i, A_i)$ makes sure than $\sigma_i(s_i, A_i)$ deems $(s_j', t_j(s_j, A_j))$ infinitely more likely than $(s_j'', t_j(s_j, A_j))$, which was to show.

So, we have shown that every $\sigma_i(s_i, A_i)$ is cautious, respects the opponents' preferences, and induces the same preference relation on S_i as $\lambda_i(s_i, A_i)$. But, as we have seen above, this completes the proof. □

References

Asheim, G.B., 2001. Proper rationalizability in lexicographic beliefs. Int. J. Game Theory 30, 453–478.

Battigalli, P., 1997. On rationalizability in extensive games. J. Econ. Theory 74, 40–61.

Blume, L.E., Brandenburger, A., Dekel, E., 1991a. Lexicographic probabilities and choice under uncertainty. Econometrica 59, 61–79.

Blume, L.E., Brandenburger, A., Dekel, E., 1991b. Lexicographic probabilities and equilibrium refinements. Econometrica 59, 81–98.

Brandenburger, A., Friedenberg, A., Keisler, H.J., 2008. Admissibility in games. Econometrica 76, 307–352.

Dekel, E., Fudenberg, D., 1990. Rational behavior with payoff uncertainty. J. Econ. Theory 52, 243–267.

Myerson, R.B., 1978. Refinements of the Nash equilibrium concept. Int. J. Game Theory 7, 73–80.

Pearce, D., 1984. Rationalizable strategic behavior and the problem of perfection. Econometrica 52, 1029–1050.

Schuhmacher, F., 1999. Proper rationalizability and backward induction. Int. J. Game Theory 28, 599–615.

Schulte, O., 2003. Iterated backward inference: An algorithm for proper rationalizability. In: Proceedings of TARK IX (Theoretical Aspects of Reasoning about Knowledge).

Stahl, D., 1995. Lexicographic rationalizability and iterated admissibility. Econ. Letters 47, 155–159.

Tan, T., Werlang, S.R.C., 1988. The Bayesian foundations of solution concepts of games. J. Econ. Theory 45, 370–391.

Zermelo, E., 1913. Über eine Anwendung der Mengenlehre auf die Theorie des Schachspiels. In: Proceedings Fifth International Congress of Mathematicians, vol. 2, pp. 501–504.