

Dynamic Coalitional Agreements - coalitional rationalizability in multi-stage games

Attila Ambrus*
Work in progress
04/28/2003

Abstract

This paper extends the concept of coalitional rationalizability of Ambrus(01) to incorporate sequential rationality in multi-stage games with observable actions and incomplete information. Agreements among players are implicit, it is assumed that players cannot communicate with each other during the game. They reflect a reasoning procedure which entails restricting strategies in a mutually advantageous way. They can be conditional on observed histories and players' types, which corresponds to allowing players to make agreements ex post and along the course of play. An agreement that is conditioned on a history is evaluated from the point of view of that history. This introduces a dynamic interaction among coalitional agreements with features of both backward and forward induction. Coalitional agreements iteratively define the set of extensive form coalitionally rationalizable strategies. This solution concept has a number of analogous properties with normal form coalitionally rationalizability. It is always nonempty. The set of outcomes consistent with it is a subset of the outcomes consistent with extensive form rationalizability, and it is robust to the order in which agreements are made. In games of perfect information extensive form coalitional rationalizability is outcome equivalent to extensive form rationalizability. Perfect coalition-proof Nash equilibria and renegotiation-proof Nash equilibria do not have to be contained in the solution set, even in two-player games, because those concepts do not imply forward induction reasoning. An alternative notion of extensive form coalitional rationalizability is also provided, assuming that coalitional agreements can only be made ex ante, but sequential individual rationality is maintained.

* Princeton University and Yale University. I thank Stephen Morris, David Pearce and especially Dilip Abreu for useful comments and suggestions. I also benefitted from conversations on the topic with Faruk Gul, Wolfgang Pesendorfer and Marciano Siniscalchi. I am grateful to the Cowles Foundation for the fellowship they granted me when I was working on this paper.

1 Introduction

Coalitional reasoning in noncooperative games means that subgroups of players with similar interest make an explicit or implicit agreement regarding what strategies to play in the game. In an equilibrium framework this reasoning leads to the concept of coalitional deviations, and to refinements of Nash equilibrium along the lines of requiring stability against coalitional deviations. The two main solution concepts proposed in the literature that address the issue of coalitional deviations are strong Nash equilibrium (see Aumann[59]) and coalition-proof Nash equilibrium (see Bernheim, Peleg and Whinston[87]). Unfortunately neither of these concepts can guarantee existence of a solution in any natural class of games, suggesting that coalitional reasoning imposed on top of Nash equilibrium might lead to contradictions. Nonexistence of strong Nash equilibrium is especially severe. Coalition-proof Nash equilibrium cannot solve the problem of existence either, despite putting restrictions on the set of allowable coalitional deviations in a manner that raises various conceptual problems (see Ambrus[01]).

Ambrus[01] proposes to deal with coalitional agreements in normal form games in a non-equilibrium framework. Instead of imposing equilibrium players are just required to be Bayesian decision makers and play best responses to conjectures that are compatible with the theory and these conjectures are not required to be correct. In this setting groups of players make implicit agreements to restrict their play to certain subsets of the strategy set. Ambrus shows that these agreements are always compatible with each other, and he proposes the solution concept of coalitional rationalizability which is obtained through an iterative procedure of agreements among subgroups of players to restrict their play to subsets of the strategy space. The set of coalitionally rationalizable strategies is always nonempty and it is a coherent set (every strategy in it is a best response to some conjecture with support inside this set, and every best response to every conjecture with support inside the set is in the set). The solution concept is a refinement of rationalizability. There is no containment relationship between the set of Nash equilibria and the set of coalition-proof Nash equilibria on one hand, and the set of coalitionally rationalizable strategies on the other hand. But every game has a Nash equilibrium that is inside the set of coalitionally rationalizable strategies. Furthermore it is shown that the solution set is insensitive to the order in which coalitional agreements are made, the same way as iterated deletion of strategies that are never best responses is insensitive to the order of deletion.

This paper addresses the issue of how coalitional rationalizability can be extended to extensive form games in a way that players are sequentially rational and foresee the coalitional agreements that are going to be made throughout the game. The assumptions that players are Bayesian decision makers and that

subgroups of players make all restrictions that are mutually advantageous for them is maintained, and sequential rationality is added as a new requirement. This introduces a dynamic element to coalitional agreements. Players can make an agreement that restrict their play at different stages of the game. Players can make an agreement at some stage of the game because they foresee that another coalitional agreement will be made at a subsequent stage of the game. Finally, players can make an agreement to jointly signal strategic intent to influence other players' choices at subsequent stages of the game.

The agreements players can make during the game are implicit in our construction, they reflect a reasoning procedure on the part of players that what restrictions on strategies are mutually advantageous at a certain stage of the game. This reasoning procedure is based on the publicly known description of the game. It is assumed that players cannot communicate to each other during the game. Communication during the game crucially changes the analysis. First, it can introduce correlation into play, while in this paper we assume that players choose their strategies independently of each other. Second, it can make players to revise their conjectures on other players' strategies even if they were not surprised by observed action choices, and therefore invalidate signaling strategic intent by action choices made before communication.¹ But we claim that communication *before the game* does not invalidate our results, in the sense that outcomes that are inconsistent with extensive form coalitional rationalizability will not be played by coalitionally rational players even in the presence of pre-play communication. There are a lot of examples though in which it is reasonable to expect players to make more restrictions on strategies played if they can communicate to each other before the game.

We restrict attention to multi-stage games with observable actions and incomplete information. Later we briefly discuss the additional issues that arise in general extensive form games. In this class of games we define the concept of history-based restrictions. A history-based restriction is an implicit agreement among a subgroup of players - a coalition - to restrict the continuation strategies that they play from that history on. History-based strategies are evaluated from the point of view of the history they are based on. This corresponds to assuming that players can make agreements along the course of play as the game progresses, or equivalently that they cannot commit themselves not to make a coalitional agreement if it becomes desirable at some stage of the game. A history-based restriction by a coalition is defined to be supported if every player in the coalition has a higher expected payoff conditional on the history being reached if players make the restriction and switch to playing only continuation strategies that are compatible with the restriction than if he plays a strategy which is outside the restriction. Based on this definition an iterative procedure is proposed, in which at every stage every coalition makes all supported

¹The introduction of public randomization devices during the game has similar effects, as shown by Gul and Pearce[96] in the context of equilibrium analysis.

history-based restrictions. The set of extensive form coalitionally rationalizable strategies is defined to be the set of strategies that survive the above procedure.

We show that extensive form coalitional rationalizability is a generalization of normal form coalitional rationalizability, and that it has similar properties to the latter. It is always nonempty. The set of outcomes that are consistent with extensive form coalitional rationalizability is always a subset of the outcomes that are consistent with extensive form rationalizability. The set of outcomes that are consistent with the solution set is insensitive to the order in which supported restrictions are made. A generalization of the concept of a coherent set is provided for a nested sequence of restrictions and it is shown that the iterative procedure defining the set of extensive form coalitionally rationalizable strategies gives a nested sequence of restrictions that is coherent. This gives an interpretation to the set of coalitionally rationalizable strategies that is similar to the interpretation that Battigalli[97] provides for the set of extensive form rationalizable strategies. It corresponds to the strategies that can be weak sequential best responses if players' conjectures at every information set are required to be consistent with the highest level of coalitional rationalizability that is consistent with the information set.

There is no containment relationship with Nash equilibrium, sequential equilibrium or trembling hand perfect equilibrium, but it is shown that every multi-stage game has a sequential equilibrium which is contained in the set of extensive form coalitionally rationalizable strategies. Perfect coalition-proof Nash equilibria, defined in Bernheim, Peleg and Whinston[87], do not have to be inside the set of extensive form coalitionally rationalizable strategies either. Unlike in normal form games, this is true even in two-player games, and in two-player finitely repeated games as well, where perfect coalition-proof Nash equilibrium coincides with renegotiation-proof Nash equilibrium (see for example Bernheim, Peleg and Whinston[87] and Benoit and Krishna[93]). This is shown to be part of a more general issue, that perfect coalition-proof and renegotiation-proof Nash equilibria might fail to be extensive form rationalizable, because they do not capture forward induction considerations that extensive form rationalizability and extensive form coalitional rationalizability do.

In games of perfect information extensive form coalitional rationalizability is shown to be outcome-equivalent to extensive form rationalizability. This establishes that in games of perfect information there are no possibilities for players to make coalitional agreements.

Extensive form coalitional rationalizability is defined for multi-stage games with incomplete information. This makes it possible to analyze coalitional agreements among players which are conditional on players' types. We allow for restrictions to be conditioned over types, or equivalently we allow for agreements

among types of players. This corresponds to assuming that just like players cannot commit not to make an agreement at a history if that agreement is mutually advantageous at the history, they cannot commit not to make an agreement that restricts the play of some of their types, if ex post the agreement is mutually advantageous from the point of the view of the types involved. We consider this assumption to be the natural extension of sequential rationality to coalitional agreements in incomplete information environments. Nevertheless we provide an alternative definition of coalitional rationalizability, ex ante coalitional rationalizability that takes the position that coalitional agreements are not credible if they are not mutually advantageous ex ante. This implies players cannot make coalitional agreements ex post and along the course of play, and every agreement is evaluated from an ex ante point of view. However the assumption of sequential (individual) rationality is maintained and shown to interact with ex ante coalitional agreements in a nontrivial way.

2 Motivating examples and an informal account of extensive form coalitional rationalizability

The examples in this section are intended to demonstrate the richness of coalitional interaction in multi-stage games with observable actions. The formal definition of this class of games, presented in the next section, requires that after every publicly observed nonterminal history every player takes an action. This is only for notational convenience though, since some players can have trivial action choices after a given history, which is equivalent to assuming that those players do not make an action choice at that history. In the examples below if a player has a trivial action choice after some history, we leave that action choice out from the description of the game.

If players can make coalitional agreements as the game progresses, conditional on a given history being reached, and we assume that players foresee these agreements, then there is a dynamic interaction among coalitional agreements. An agreement that is made at some history might trigger another agreement at an earlier stage of the game, as the game of Figure 1 below demonstrates. This is an analogue, in the context of coalitional agreements, of the backward induction logic implied by sequential rationality. Furthermore, players might make a coalitional agreement to jointly signal strategic intent at some later stage of the game, as the game of Figure 2 shows. Therefore coalitional agreements to restrict strategies in extensive form games are intertwined with the type of forward induction logic implied by extensive form rationalizability.




Figure 1

In the game of Figure 1, if the information sets of player 3 and player 4 are reached, then (L, l) yields the best payoff that they can get, so it is mutually advantageous for them to make an implicit agreement not to play R and r . Knowing this, it is mutually advantageous for players 1 and 2 to make an implicit agreement not to play R and r . This is along the lines of backward induction or sequential logic, combined with coalitional agreements. Note that coalitional rationalizability in the normal form of this game does not eliminate any strategies. In particular playing (L, l) is not a supported restriction for players 3 and 4, because they might think that player 1 plays R or player 2 plays r with probability 1, in which case it does not matter what strategies players 3 and 4 choose. But then (L, l) is not a supported restriction for players 1 and 2, since they cannot be sure that players 3 and 4 play (L, l) . Finally, the coalition of all players does not have a supported restriction, because the favorite outcome of players 3 and 4 is not (L, l, L, l) , but (R, r, R, r) . Even applying perfect coalitional rationalizability (see Ambrus[01]) to the normal form of the game does not eliminate any strategies. In finite games perfect coalitional rationalizability is equivalent to one round of elimination of weakly dominated strategies and taking the set of coalitionally rationalizable strategies of the remaining set. Since no strategy is weakly dominated in the normal form representation, the previous argument establishes that perfect coalitional rationalizability does not put any restriction on strategies that can be played.




Figure 2

In the game of Figure 2, it is undominated for player 1 to choose I only if he plans to play L afterwards. Similarly for player 2. Knowing this, player 3 should play L after observing (I, I) . Therefore if play starts with (I, I) , the only reasonable outcome is the one belonging to the endnode on the left, giving a payoff vector $(3, 3, 1)$ which is the best possible payoff to players 1 and 2. Coalitional reasoning should then induce players 1 and 2 to start out by playing I . Note that extensive form rationalizability does not eliminate other strategies by player 1 and player 2, because O is a best response for each of them if they allocate a high enough probability to the other choosing O . Some form of coalitional reasoning is needed to get around the coordination problem. Also note that coalitional rationality in the normal form of the game is ineffective in deleting strategies, because it cannot reproduce the forward induction argument needed in the first place. And since player 3's favorite outcome is when players 1 and 2 play (O, O) , he cannot be part of a mutually advantageous agreement. We call the above reasoning procedure "coalitional forward induction", because

players 1 and 2 make a coalitional agreement that involves signaling their joint strategic intent to player 3.

Another implication of players using sequential logic in making coalitional agreements is that in games of incomplete information players evaluate coalitional agreements conditional on their realized types. If an agreement is advantageous for a realized type of some player, then it is advantageous for him to make the agreement even if the same agreement is not advantageous for some other type of the same player. Consider the game of Figure 3.




Figure 3

In the above game player 1 has two possible types, type A and type B, which are realized with equal probability. Consider now the implicit agreement between players 1 and 2 to play strategies R and r . This is definitely a desirable agreement from the point of view of type B of player 1, since it yields him the highest possible payoff in the game, but not from the point of view of type A of player 1. Note however that the agreement is a desirable one from the point of view of player 2 even if he assumes that only type B of player 1 follows the agreement. If type B of player 1 plays R and he plays r , then he gets an expected payoff of at least $3/2$, while playing l can never yield an expected payoff of higher than 1. Therefore it is in the interest of player 2 to make an implicit agreement with type B of player 1 to play (R, r) . Getting back to type B of player 1, although the agreement is advantageous for him, he knows that it is not advantageous for type A. This, however does not mean that it is not in his interest to make the agreement to play (R, r) . If nature selected type B

for player 1, then player 1, when called upon to choose an action knows that type B was realized and not type A. Therefore sequential logic suggests that he evaluates the agreement from type B's point of view, ignoring whether the agreement would have been advantageous for type A had he been realized.

The formal construction of extensive form coalitional rationalizability, presented in the next section, allows players to make coalitional agreements conditional on realized types and conditional on histories being reached. Since agreements are conditional on realized types, formally we allow types of players to evaluate and make restrictions. This can be interpreted literally as well, that it is the types of players who evaluate and possibly make coalitional agreements with each other. In most of the paper our terminology reflects this view and in verbal analysis we will talk about "agreements among types". This is only for ease of exposition though, as our view is that these agreements represent agreements among players conditional on realized types. If a player makes an implicit agreement with some, but not all types of another player, then just like in the example of Figure 3, the agreement is only advantageous for him if it is advantageous for him even if the types who are not part of the agreement play some other strategies than those specified by the agreement. Since this evaluation is done using expected payoffs, the more likely a player considers that he is facing the types who are involved in the agreement, the more likely it is that the agreement is advantageous for him. This introduces an interesting interplay with allowing the agreements to be conditioned on histories being reached. If at a given history a player thinks that it is very likely that he faces certain types of the other players, then an agreement with those types conditional on the history can be advantageous even if at the beginning of the game he did not think that those types were very likely. Consider the game of Figure 4.




Figure 4

In this game player 1 has three possible types: A, B and C, but for ease of exposition nature's move is not made explicit in the figure. At every final outcome, the first payoff belongs to type A of player 1, the second to type B of player 1, the third to type C of player 1 and the fourth to player 2. Assume that nature chooses each type of player 1 with probability $1/3$. Note that playing r at the beginning of the game is a dominant action for type C of player 1. On the other hand, playing l at the beginning of the game is a dominant action for type A of player 1. Therefore after history l was observed, player 2 should allocate at least probability $1/2$ to facing type A of player 1. He should allocate probability 0 at that stage to type C of player 1 being realized, and some probability between $1/2$ and 1 to type A being realized, depending on what he thinks about the action choice at the beginning of the game of type B of player 1. But then playing (L, l) after history l is an advantageous agreement for player 2 and type A of player 1. It is advantageous for type A of player 1 because it yields the highest possible payoff for him in the game. And it is advantageous for player 2 because given that history l was reached the agreement guarantees him an expected payoff of at least $3/2$, while playing r could yield at most a payoff of 1. Note that the agreement is supported only because conditional on l being reached the probability of type A of player 1 being realized is higher than $1/3$, the prior probability of type A.

The next example points out that agreements that are advantageous conditional on a history might trigger another agreement that causes that history never being reached, just like sequential rationality reasoning applied to continuation strategies from a history can ultimately rule out that history being reached.




Figure 5

In the game of Figure 5, conditional on the subgame consisting of player 2's and player 3's information set is reached, (L, l) gives a higher payoff to both of them than what any other strategies yield. Therefore if the subgame is reached, it is mutually advantageous for them to coordinate play accordingly. This is not changed by the fact that player 1, anticipating the above agreement is better off playing O , which gives the worst payoff to players 2 and 3. Players 2 and 3 would like player 1 to believe that they will not coordinate on (L, l) , but if their subgame is reached they have every incentive to play those strategies.

In the previous game an agreement conditional on a history ultimately hurts the players because it causes another player to restrict his play at an earlier stage of the game in a way that is not advantageous for the players making the first agreement. The next example demonstrates a different conflict of interest between a player's interests at different points of a game, in games of incomplete information.




Figure 6

The game of Figure 6 has the same structure as the game of Figure 3, and similarly (R, r) is an advantageous agreement for player 2 and type B of player 1. But note that in this game the agreement is not only disadvantageous from the point of view of type A of player 1, but from an ex ante point of view as well. Ex ante player 1 is clearly worse off if player 2 plays R (then player 1's ex ante expected payoff is at most $-7/2$) than if player 2 plays L (a guaranteed nonnegative ex ante expected payoff). This means that if ex ante player 1 could commit not to make the agreement (R, r) then he would be better off doing that. But if he cannot make such a commitment, then since for type B of player 1 it is still unambiguously advantageous to make the agreement and therefore he is willing to make it. This makes type A so much worse off that offsets the positive payoff effect for type B from an ex ante point of view. This reflects a conflict of interest between a player's interest ex ante and her ex post interest after a certain type realization.

In the next section we formally define a solution concept, the set of extensive form coalitionally rationalizable strategies, that is consistent with the considerations in the above examples. This section concludes with an informal description of the construction of extensive form coalitional rationalizability.

The key element in defining extensive form coalitional rationalizability is defining the coalitional agreements in extensive form games that are unambiguously in the interest of the participants. Since the type of agreements we

consider are those in which participants agree upon restricting their play to a certain subset of the strategy space, we call these agreements “supported restrictions.” The definition of supported restrictions in multi-stage games has to confront two issues not present in the normal form. First, the game has a sequential structure, so making an agreement can become strictly in the interest of the members of a coalition at a certain stage of the game, even if it was not necessarily so at the beginning of the game. Second, in multi-stage games with incomplete information the realized type of a player is private information, and the interest of different types of the same player might be very different. In the formal construction of the next section these two issues are taken up simultaneously and the definitions allow for both multiple stages and incomplete information. But here we separate the two issues and first provide an informal description of supported restrictions in multi-stage games with complete information.²

The definition of supported restriction has to specify for any proposed restriction that from what point of view the players involved evaluate the restriction. Since in the class of games we examine past actions are observable, the natural points of evaluation (and the one compatible with the usual concept of sequential rationality) are the histories which have the feature that the restriction is essentially on continuation strategies from that history. In particular every restriction can be evaluated from the null history, but to be evaluated from any other history, the restriction has to satisfy the above requirement. By evaluation from the point of view of a certain history we mean that players who contemplate the restriction compare updated expected payoffs conditional on that history being reached.

A restriction is unambiguously advantageous for a player if his updated expected payoff at the history from which the restriction is evaluated is strictly lower if the agreement is not made and it is optimal for him to play a continuation strategy that is outside the restriction, then if the agreement is made and every other player in the coalition only plays strategies that are inside the restriction. This implies that only those scenarios are considered in the payoff comparison in which it is optimal for the player to play a strategy outside, and in all these scenarios the player should be strictly better off by switching to making the agreement once the above information set is reached. By the strategy being optimal we mean that it is a weak sequential best response to a consistent conjecture (for the formal definition of these concepts see the next section).

Just like in normal form games, the payoff comparison is done by fixing the marginal conjecture concerning players outside the coalition, but switching to making the agreement is required to be strictly in the interest of the player for

²Section 5 provides a more formal analysis both in games of sequential nature but no private information (multi-stage games with complete information) and games with private information but no sequentiality (normal form Bayesian games).

every fixed marginal that is compatible with the optimality of playing a strategy outside the restriction. This corresponds to the assumption that players outside the coalition cannot observe whether an agreement is made or not and therefore cannot condition their strategies on that event.

A restriction is supported if it is unambiguously advantageous for all players involved.

Given this definition of a supported restriction the iterative procedure that defines the set of extensive form coalitionally rationalizable strategies is similar to the corresponding iterative procedure that defines the set of coalitionally rationalizable strategies in normal form games. Players first consider the set of all strategies. Then it is assumed that all restrictions that are supported from the point of view of some history are made. In the next step, players' conjectures are assumed to be concentrated on strategies that are not eliminated in the first step. Given that assumption, again all restrictions that are supported from some history are made. The procedure continues in this manner until it reaches a set from which there are no history-based restrictions. That set, which is shown to be always nonempty, is called the set of extensive form coalitionally rationalizable strategies.

Incomplete information is incorporated in the framework essentially by allowing different types of the same player to make different agreements. Essentially the set of players is extended to be the set of possible types. Section 5 shows that this analogy is exact in Bayesian normal form games. The set of extensive form coalitionally rationalizable strategies in these games is the same as the set of coalitionally rationalizable strategies in the normal form game that is obtained from the Bayesian game by treating types as separate players and defining their payoffs using expected payoffs of the original game. In general, the definition of supported restriction allows for the possibility of agreements among general coalition of types, including coalitions which involve some but not all types of a given player. But keeping the central feature of the definition from the complete information context, such a restriction is only supported if it is unambiguously in the interest of all types involved in the restriction, even in cases when the types outside the coalition play strategies that are not compatible with the restriction.

3 Construction and basic properties

We consider finite multi-stage game of incomplete information with observed actions. At the beginning of the game chance makes a move, selecting the type of each player. Then the realized types play a multi-stage game, where

at each nonterminal stage players observe all previous action choices and then simultaneously choose actions. At any stage the set of feasible actions for a player is allowed to be a singleton, in which case that player effectively does not choose an action at that stage.³

A multi-stage game of incomplete information with observed actions G is formally defined as a structure $\langle N, (\Upsilon_i)_{i \in N}, (\varphi_i)_{i \in N}, H, (u_i)_{i \in N} \rangle$, where N is the set of players, Υ_i is the set of possible types of player i , φ_i is player i 's belief, conditional on his type, over the types of other players, H is the set of feasible public histories and u_i is the payoff function of player i over terminal histories, conditional on his type. We define these concepts formally below.⁴

For every $i \in N$, φ_i is a function $\Upsilon_i \rightarrow \Delta(\prod_{j \in N/\{i}} \Upsilon_j)$, where $\Delta(\prod_{j \in N/\{i}} \Upsilon_j)$ is the set of probability distributions over $(\prod_{j \in N/\{i}} \Upsilon_j)$. $\varphi_i(\tau_i)$ represents the belief of type τ_i over the distribution of types of the other players. We do not assume in the construction that this function is derived from a prior common probability distribution on Υ , although we make this assumption in most concrete examples we present in the paper.⁵ Furthermore, even if the functions φ_i can be derived from a common prior probability distribution φ on Υ , φ does not have to be a product measure (types of different players do not have to be independent).

H is the set of possible public histories. It consists of the null history \emptyset and sequences of the form $(a_1^1, \dots, a_N^1, a_1^2, \dots, a_N^2, \dots, a_1^k, \dots, a_N^k) \equiv (a^1, \dots, a^k)$ for some $k \geq 1$. The member of the sequence a_i^m (for $1 \leq m \leq k$, $i \in N$) represents the action chosen at stage m by player i . History $h = (a^1, \dots, a^k)$ is *terminal* if there is no a^{k+1} such that $(a^1, \dots, a^k, a^{k+1}) \in H$. Let H^Z denote the set of terminal histories. It is assumed that if $k \geq 2$ and $(a^1, \dots, a^k) \in H$, then also $(a^1, \dots, a^{k-1}) \in H$. Furthermore, for every $i \in N$ if $(a^1, \dots, a^{k-1}, a^k) \in H$ and $(a^1, \dots, a^{k-1}, b^k) \in H$, then $(a^1, \dots, a^{k-1}, a_1^k, \dots, a_{i-1}^k, b_i^k, a_{i+1}^k, \dots, a_N^k) \in H$.

For every $i \in N$, u_i is a function $\Upsilon_i \times H^Z \rightarrow R$. For every $\tau_i \in \Upsilon_i$ and $h \in H^Z$, $u_i(\tau_i, h)$ represents the payoff of type τ_i of player i at terminal history h .

We assume that N , $\Upsilon \equiv \prod_{i \in N} \Upsilon_i$ and H are finite.

³changing the formal definition of the class of games to explicitly allow for the possibility that only a subset of players make action choices after some history is inconsequential for the analysis presented below, but it would make the notation more cumbersome.

⁴for a more elaborate construction of extensive form games and multi-stage games in particular see for example the textbook of Osborne and Rubinstein[94].

⁵Starting from some prior common prior $\varphi \in \Delta(\Upsilon)$ which has the property that every type of every player is realized with positive probability ($\forall i \in N$, $\tau_i \in \Upsilon_i$ it holds that there is $v \in \Upsilon$ such that $v_i = \tau_i$ and $\varphi(v) > 0$) it is straightforward to construct the posterior beliefs of a type on other players' types: for every $i \in N$, $\tau_i \in \Upsilon_i$ and $\varsigma_{-i} \in \prod_{j \in N/\{i}} \Upsilon_j$

$$\varphi_i(\tau_i)(\varsigma_{-i}) = \varphi(\varsigma_{-i}, \tau_i) / \sum_{v: v_j = \tau_j} \varphi(v).$$

Action sets and strategies

The above definition of histories implies that at every $h \in H/H^Z$ there is a nonempty set of feasible actions $A_i(h)$ for every $i \in N$, such that $(h, a_1, \dots, a_N) \in H$ iff $a_i \in A_i(h) \forall i \in N$. The construction implies that the set of feasible actions after some history is the same for all types of the same player. This assumption is not crucial though and made primarily for notational convenience. Furthermore, a game in which different types of the same player have different action sets can be incorporated in this framework for the purposes of our analysis by defining payoffs for infeasible actions such that choosing those actions are always strictly dominated for the corresponding types.

A pure strategy of player $i \in N$ is a function that allocates an element of $A_i(h)$ to every $h \in H/H^Z$. Let S_i be the set of pure strategies of player i and let $S = \times_{i \in N} S_i$. Every $s \in S$ is associated with a unique terminal history $h(s) \in H^Z$ and therefore $u_{\tau_i}(s) = u_i(h(s), \tau_i)$, the payoff that type τ_i of player i gets if the realized types play strategy s is a well-defined function.

Extended strategy profiles

Below we assume that player types can make coalitional agreements and therefore we extend the set of strategies from the set of players to the set of types. An extended strategy profile specifies a strategy for every type of every player. Let \underline{S} be the set of strategy profiles: $\underline{S} = \times_{i \in N} S_i^{\Upsilon_i}$. \underline{S} is just the product of players' strategy sets, where the number of components that are equal to a particular player's strategy set is the number of types the given player has. Formally, $\underline{S}_{\tau_i} = \underline{S}_{\tau'_i} = S_i \forall i \in N$ and $\tau_i, \tau'_i \in \Upsilon_i$. In the construction below coalitional agreements make restrictions on \underline{S} such that the resulting $\underline{A} \subset \underline{S}$ typically does not have the property that $\underline{A}_{\tau_i} = \underline{A}_{\tau'_i} \forall i \in N$ and $\tau_i, \tau'_i \in \Upsilon_i$.

From now on we will use underlying to indicate that a profile or a set of profiles belongs to \underline{S} . In particular we denote a typical element of \underline{S} by \underline{s} and a typical subset of \underline{S} by \underline{A} . Furthermore, define $\underline{S}_{-i} \equiv \times_{j \in N/\{i\}} \times_{\tau_j \in \Upsilon_j} \underline{S}_{\tau_j}$. Similarly for $\underline{s} \in \underline{S}$ let $\underline{s}_{-i} = \times_{j \in N/\{i\}} \times_{\tau_j \in \Upsilon_j} \underline{s}_{\tau_j}$ and for $\underline{A} = \times_{j \in N} \times_{\tau_j \in \Upsilon_j} \underline{A}_{\tau_j} \subset \underline{S}$ let $\underline{A}_{-i} = \times_{j \in N/\{i\}} \times_{\tau_j \in \Upsilon_j} \underline{A}_{\tau_j}$.

Beliefs and belief processes concerning strategies

We treat players' beliefs concerning the distribution of other players' types and concerning the strategies played by those types separately. Above we defined beliefs over the distribution of other players' types as part of the description of the game.

The set of beliefs that at a given point of the game a type of player i can have over the strategy choices of other players types is the set of probability distributions over \underline{S}_{-i} . Let $\Omega_{-i} = \Delta(\underline{S}_{-i})$.

Note that we allow for correlated conjectures over strategies played by the others. The qualitative results of the paper would remain unchanged though if we restricted conjectures to be independent.

Let Θ_{-i} be the set of functions mapping H/H^Z into Ω_{-i} . We call Θ_{-i} the belief processes of player i .⁶ A belief process θ_{-i} specifies a belief for player i concerning other players' strategies after every history. Let $\theta_{-i}[h]$ denote the belief that θ_{-i} specifies after $h \in H/H^Z$.

A belief process of player i specifies conjectures only on strategies of *other* players' types and therefore Θ_{-i} is the set of belief processes for *every type* of player i . On the other hand we do not require the conjectures of different types of the same player to be necessarily the same. The latter assumption, often made in the literature on signaling games, would put extra restrictions on strategy profiles that can be expected to be played in the game and would lead to a stronger solution concept than that presented below. In some contexts, in particular if players' beliefs over types come from a prior probability distribution over types that is independent (and therefore different types of the same player have the same beliefs over the other players' types) this extra assumption is more appealing than in others. For a more detailed discussion of this issue, see Sobel, Stole and Zapater[90] and Battigalli and Siniscalchi[01].

Note that although we allow for correlated conjectures concerning strategies of other players, the framework above does not allow correlation between a player's belief concerning other players' types and the conjecture he has concerning other players' strategies. Our model could be extended though to allow for this kind of correlation as well. A related generalization of our framework would be to define belief processes over both types and strategies of other players⁷ and therefore allow players to change their beliefs concerning the distribution of other players' types, not just concerning their strategies, after surprise events. In the model we present the belief of a player type over the distribution of other players' types is fixed throughout the game. This does not imply that after any history a player type's belief over what types of other players he is facing in the game is constant, only that the differences are obtained solely by Bayesian updating based on the belief process over strategies of other players' types. Players do make inferences on what types of other players were realized, but they do not change their belief concerning the a priori distribution of types.

Similarly to the above, imposing the stronger requirement of "strategic independence" (see Battigalli[96]), on top of independent conjectures, does not change the qualitative results of the paper.

⁶elsewhere in the literature these functions are called updating systems or conditional (probability) systems.

⁷for a construction of belief processes along these lines, see Battigalli and Siniscalchi[01].

Expected payoff

Let $u_{\tau_i}(s_i, \omega_{-i})$ be the expected payoff of type τ_i of player i if he plays strategy s_i and has belief $\omega_{-i} \in \Omega_{-i}$:

$$u_{\tau_i}(s_i, \omega_{-i}) = \sum_{v_{-i} \in \Upsilon_{-i}} \left[\sum_{s_{-i} \in S_{-i}} \sum_{\underline{t}: t(v_{-i})=s_{-i}} u_{\tau_i}(s_i, s_{-i}) \cdot \omega_{-i}(\underline{t}) \right] \cdot \varphi(\tau_i)(v_{-i}).$$

A history being reached

We say that $h' = (b^1, \dots, b^m) \in H$ is a predecessor of $h = (a^1, \dots, a^k) \in H$ if $k < m$ and $b^l = a^l \forall l = 1, \dots, m$. Let $\text{pred}(h)$ be the set of predecessors of any $h \in H$. We say that $h' \in H$ is a successor of $h \in H$ if h is a predecessor of h' . Let $\text{succ}(h)$ be the set of successors of any $h \in H$. We say that $h' \in H$ is an immediate predecessor of $h \in H$ if h' is a predecessor of h and there is no $h'' \in H$ such that h'' is a predecessor of h a successor of h' . Let $\text{imp}(h)$ be the immediate predecessor of any $h \in H/\{\emptyset\}$.

Next we formally define that a strategy and a strategy profile reaches a given history, and that a belief reaches a given history for a player type. These concepts are needed for the construction for two reasons. First, to define consistency of belief processes. Second, to define formally whether a restriction of strategies is supported by a player type. This is because we assume that when a player type evaluates a coalitional agreement on restricting strategies conditional on a history, he only considers scenarios that make it possible that play reaches that history.

Let $i \in N$, $h = (a^1, \dots, a^k)$ and $h' = (a^1, \dots, a^m) \in \text{pred}(h)$. Then $a_i(h, h') \equiv a_i^{m+1}$, the action that h specifies for player i after the sub-history h' . We say that $s_i \in S_i$ reaches h if $s_i(h') = a_i(h, h') \forall h' \in \text{pred}(h)$. Similarly, $s \in S$ reaches h if $s_i(h') = a_i(h, h') \forall h' \in \text{pred}(h)$ and $i \in N$. We say that $\omega_{-i} \in \Omega_{-i}$ reaches h for type $\tau_i \in \Upsilon_i$ if there is $v_{-i} \in \Upsilon_{-i}$ and $\underline{s}_{-i} \in \underline{S}_{-i}$ such that $\varphi(\tau_i)(v_{-i}) > 0$, $\omega_{-i}(\underline{s}_{-i}) > 0$ and $\underline{s}_{v_j}(h') = a_j(h, h') \forall j \in N/\{i\}$.

Consistent belief processes

We only consider belief processes that satisfy two consistency properties. If a player's belief after some history reaches a subsequent history, then after the latter history the player shouldn't revise his conjecture. And after any history a player should have a belief that reaches that history.

Definition: a belief process $\theta_{-i} \in \Theta_{-i}$ is *consistent* if

- (i) if $h, h' \in H/H^Z$, $h' \in \text{pred}(h)$ and $\theta_{-i}(h')$ reaches h , then $\theta_{-i}(h) = \theta_{-i}(h')$
- (ii) $\forall h \in H/H^Z$, $\theta_{-i}(h)$ reaches h .

Let Θ_{-i}^c be the consistent belief processes of player i .

Belief processes concentrated on a subset of strategies and sequences of nested restrictions

In what follows we assume that players make a series of agreements on restricting their set of strategies, resulting in a nested sequence of subsets of the original strategy space. At each stage of this procedure we assume that players believe that all previously agreed upon restrictions are indeed made. Below it is formally defined that a player believes that a restriction or a sequence of restrictions is made, in terms of conditions on his belief process.

Let $\underline{A} = \times_{i \in N} (\times_{\tau_i \in \Upsilon_i} \underline{A}_{\tau_i})$ be such that $\underline{A}_{\tau_i} \subset S_i$ and $\underline{A}_{\tau_i} \neq \emptyset \forall i \in N$ and $\tau_i \in \Upsilon_i$. Then for every $i \in N$ let $\Omega_{-i}(\underline{A}) \equiv \{\omega_{-i} \in \Omega_{-i} \mid \omega_{-i}(\underline{s}_{-i}) = 0 \forall \underline{s}_{-i} \notin \underline{A}_{-i}\}$. We call $\Omega_{-i}(\underline{A})$ the set of beliefs of player i that are concentrated on \underline{A} . Similarly let $\Theta_{-i}^c(\underline{A}) \equiv \{\theta_{-i} \in \Theta_{-i}^c \mid \theta_{-i}(h) \in \Omega_{-i}(\underline{A}) \forall h \in H/H^Z \text{ such that } \underline{A}_{-i} \text{ reaches } h\}$. We call $\Theta_{-i}^c(\underline{A})$ the set of consistent belief processes of player i that are concentrated on \underline{A} .

Let now $k \geq 1$ and $(\underline{B}^1, \dots, \underline{B}^k)$ be such that for every $m = 1, \dots, k$ it holds that $\underline{B}^m = \times_{i \in N} (\times_{\tau_i \in \Upsilon_i} \underline{B}_{\tau_i}^m)$ where $\underline{B}_{\tau_i}^m \subset S_i$ and $\underline{B}_{\tau_i}^m \neq \emptyset \forall i \in N$ and $\tau_i \in \Upsilon_i$.

Assume furthermore that $\underline{B}^m \supset \underline{B}^{m+1} \forall m = 1, \dots, k-1$. We call sequences of strategy sets like that sequences of nested restrictions. Then $\Theta_{-i}^c(\underline{B}^1, \dots, \underline{B}^k) \equiv \bigcap_{m=1, \dots, k} \Theta_{-i}^c(\underline{B}^m)$. The belief systems in $\Theta_{-i}^c(\underline{B}^1, \dots, \underline{B}^k)$ are such that at histories reached by \underline{B}_{-i}^k the belief specified at the history is concentrated on \underline{B}_{-i}^k , at histories not reached by \underline{B}_{-i}^k but reached by \underline{B}_{-i}^{k-1} the belief specified at the history is concentrated on \underline{B}_{-i}^{k-1} and so on, so that at any history the belief specified at the history is concentrated on the last member of the sequence $(\underline{B}^1, \dots, \underline{B}^k)$ that is compatible with the history.

Updated payoff expectations after histories

Let $h \in H/H^Z$ and $\theta_{-i} \in \Theta_{-i}^c$. Let $S_{-i}[h]$ be the set of profiles in S_{-i} that reach h . Then $p_{\tau_i}^h(s_{-i} \mid \theta_{-i})$ is the probability that type τ_i of player i , having belief process θ_{-i} , allocates to other players playing profile s_{-i} after observing history h : $p_{\tau_i}^h(s_{-i} \mid \theta_{-i}) = 0$ if $s_{-i} \notin S_{-i}[h]$ and $p_{\tau_i}^h(s_{-i} \mid \theta_{-i}) = \frac{q^h(s_{-i} \mid \theta_{-i})}{\sum_{s_{-i} \in S_{-i}(h)} q^h(s_{-i} \mid \theta_{-i})}$ if $s_{-i} \in S_{-i}[h]$, where $q^h(s_{-i} \mid \theta_{-i}) \equiv \sum_{v_{-i} \in \Upsilon_{-i}} \theta_{-i}[h](\underline{s}) \cdot \varphi(\tau_i)(v_{-i})$.

Let $s_i \in S_i$, $\tau_i \in \Upsilon_i$, $h \in H/H^Z$ and $\theta_{-i} \in \Theta_{-i}^c$. Assume h is reached by s_i . Then the updated expected payoff of type τ_i of player i , having belief process θ_{-i} and playing strategy s_i , after observing history h is $u_{\tau_i}^h(s_i, \theta_{-i}) \equiv \sum_{s_{-i} \in S_{-i}} u_{\tau_i}(s_i, s_{-i}) \cdot p_{\tau_i}^h(s_{-i} \mid \theta_{-i})$.

Replacement strategies and sequential best responses

Let $h \in H/H^Z$. We call $s_i \in S_i$ to be an h -replacement of $t_i \in S_i$ if $s_i(h') = t_i(h')$ for every $h' \in H/H^Z$ such that $h' \notin h \cup \text{succ}(h)$.

Then $s_i \in S_i$ is a best response for $\tau_i \in \Upsilon_i$ to $\theta_{-i} \in \Theta_{-i}^c$ among h -replacements if $u_{\tau_i}^h(s_i, \theta_{-i}) \geq u_{\tau_i}^h(t_i, \theta_{-i})$ for every $t_i \in S_i$ such that t_i is a h -replacement of s_i .

Let $\theta_{-i} \in \Theta_{-i}^c$. Strategy $s_i \in S_i$ is a weak sequential best response for $\tau_i \in \Upsilon_i$ to θ_{-i} if s_i is a best response for τ_i to θ_{-i} among h -replacements for every $h \in H/H^Z$ that is reached by s_i .

A weak sequential best response (from now on just best response for ease of exposition) of a player type to a belief system is a strategy that is a best response among replacement strategies after every history not excluded by the strategy.

Let $BR_{\tau_i}(\theta_{-i})$ denote the set of strategies $s_i \in S_i$ which are best responses for τ_i to θ_{-i} .

It is straightforward to establish that there always exists a best response against a consistent belief process.

History-based restrictions

A history-based restriction by a group of types represents an implicit agreement among these types that restricts the set of continuation strategies from a certain history. First we define the criterion that a restriction is based on a given history. The construction below assumes that a restriction that is based on some history is evaluated from the point of view of that history.

We say that Φ is a *coalition of types* if $\Phi = \bigcup_{i \in N} \Phi_i \neq \emptyset$, where $\Phi_i \subset \Upsilon_i \forall i \in N$.

Let $\underline{A} = \times_{i \in N} (\times_{\tau_i \in \Upsilon_i} \underline{A}_{\tau_i})$ be such that $\underline{A}_{\tau_i} \subset S_i$ and $\underline{A}_{\tau_i} \neq \emptyset \forall i \in N$ and $\tau_i \in \Upsilon_i$. Let $h \in H/H^Z$. Let Φ be a coalition of types.

Definition: $\underline{B} = \times_{i \in N} (\times_{\tau_i \in \Upsilon_i} \underline{B}_{\tau_i})$ is a *h -based restriction* by Φ given \underline{A} if $\underline{B} \subset \underline{A}$, $\underline{B} \neq \emptyset$ and

- (i) $\underline{B}_{\tau_i} = \underline{A}_{\tau_i} \forall \tau_i \notin \Phi$
- (ii) $\forall \tau_i \in \Phi$ and $a_i \in \underline{A}_{\tau_i}/\underline{B}_{\tau_i}$ it holds that a_i reaches h and $\exists b_i \in \underline{B}_{\tau_i}$ such that b_i is a h -replacement of a_i .

Condition (i) requires that a restriction by a coalition of types Φ should only restrict strategies played by those types. Condition (ii) requires that for every type in the coalition and every strategy that the restriction rules out for this type, the strategy reaches the history and that it has a replacement strategy from that history that is not ruled out by the restriction. The motivation behind this requirement is that an agreement should only be evaluated from a history h only if it is an agreement concerning continuation strategies from h . Otherwise the restriction affects what action choices can be chosen at histories that are outside $h \cup \text{succ}(h)$ and therefore it is incorrect to evaluate the agreement from h . Note that any nonempty product subset can be a *null-history based* restriction, so the definition does not rule out any possible restriction from the starting set. But in order to be labeled as an h -based restriction for some history h different

than the null history, the restriction has to satisfy the property that it only restricts the set of continuation strategies from that history.

Definition: \underline{B} is a *restriction* by Φ given \underline{A} if it is a h -based restriction for some $h \in H/H^Z$.

Outcomes compatible with a restriction

Let $\underline{A} = \times_{i \in N} (\times_{\tau_i \in \Upsilon_i} \underline{A}_{\tau_i})$ be such that $\underline{A}_{\tau_i} \subset S_i$ and $\underline{A}_{\tau_i} \neq \emptyset \forall i \in N$ and $\tau_i \in \Upsilon_i$. Then for any profile of types $\tau \in \Upsilon$ the outcomes reached by \underline{A} for type profile τ , denoted by $O_\tau(\underline{A})$ is the set of final histories that are reached by strategy profiles in $\times_{i \in N} \underline{A}_{\tau_i}$: $O_\tau(\underline{A}) = \{h \in H^Z \mid \exists s_{\tau_1} \in \underline{A}_{\tau_1}, \dots, s_{\tau_N} \in \underline{A}_{\tau_N} \text{ such that } (s_{\tau_1}, \dots, s_{\tau_N}) \text{ reaches } h\}$. If players play strategies in \underline{A} and τ is the profile of realized types, then $O_\tau(\underline{A})$ is the set of final histories that play can reach.

Sequences of restrictions that are closed under rational behavior

Let $(\underline{B}^1, \dots, \underline{B}^k)$ be a nested sequence of restrictions. We say $(\underline{B}^1, \dots, \underline{B}^k)$ is closed under rational behavior if for every $i \in N$ and $\tau_i \in \Upsilon_i$ it holds that $\theta_{-i} \in \Theta_{-i}^c(\underline{B}^1, \dots, \underline{B}^k)$ and $s_i \in BR_{\tau_i}(\theta_{-i})$ implies $s_i \in \underline{B}_{\tau_i}^k$.

A nested sequence of restrictions is closed under rational behavior if for every type of every player it holds that for every consistent belief system that is concentrated on the sequence, all best responses to this belief system are inside the last set in the sequence (and therefore in every set of the sequence).

The above defined notion is a generalization of the well-known concept of sets closed under rational behavior to extensive form games. Consider the normal form game $G = (N, (S_i)_{i \in N}, (u_i)_{i \in N})$, where N is the set of players, S_i is the set of strategies of player i and u_i is the payoff function of player i . Then $B = \times_{i \in N} B_i$ is closed under rational behavior if for every $i \in N$ the set of best responses to any conjecture of player i with support in B_{-i} is contained in B_i . In the normal form there is no need to consider nested sequences of restrictions, since the trivial information structure implies that a conjecture is concentrated on a sequence of restrictions if and only if it is concentrated on the last member of the sequence. In the extensive form one has to keep track of previous restrictions in order to establish parallels between a sequence of restrictions on beliefs and a corresponding sequence of restrictions on players' conjectures.⁸

Let \mathcal{M} denote the nested sequences of restrictions that are closed under rational behavior.

Supported restrictions

Next the definition of a supported restriction by a coalition is provided. Supported restriction formalizes the concept of an implicit agreement among a

⁸See Battigalli[97] for a more detailed discussion.

group of players that is unambiguously in the interest of every player in that group once an information set is reached.

Let $\underline{A} = \prod_{i \in N} (\prod_{\tau_i \in \Upsilon_i} \underline{A}_{\tau_i})$ be such that $\underline{A}_{\tau_i} \subset S_i$ and $\underline{A}_{\tau_i} \neq \emptyset \forall i \in N$ and $\tau_i \in \Upsilon_i$. Let Φ be a coalition of types.

Definition: \underline{B} is an *h-based supported restriction* by Φ given \underline{A} if

- (i) \underline{B} is an *h-based restriction* by Φ given \underline{A}
- (ii) $\forall \tau_i \in \Phi$, $s_i \in \underline{A}_{\tau_i} / \underline{B}_{\tau_i}$ and $\theta_{-i} \in \Theta_{-i}^c(\underline{A})$ such that $s_i \in BR_{\tau_i}(\theta_{-i})$ it holds that $u_{\tau_i}^h(s_i, \theta_{-i}) < u_{\tau_i}^h(t_i, \psi_{-i})$
 $\forall t_i, \psi_{-i}$ such that $\psi_{-i} \in \Theta_{-i}^c(\underline{B})$, $\theta_{-\Phi}(h) = \psi_{-\Phi}(h)$, $t_i \in \underline{A}_{\tau_i}$ and t_i is a best

response to ψ_{-i} among h' -replacement strategies of t_i for every $h' \in h \cup \text{succ}(h)$ that is reached by t_i .

Definition: \underline{B} is a *supported restriction* by Φ given \underline{A} if it is a *h-based supported restriction* for some $h \in H/H^Z$. A supported restriction \underline{B} by Φ given \underline{A} is *nontrivial* if $\underline{B} \neq \underline{A}$.

A *h-based supported restriction* has the property that the updated expected payoff of any type at h if he plans to play a continuation strategy that is ruled out by the restriction is strictly lower than his updated expected payoff at h if he has a belief process that is concentrated on the restriction and plays a best response to that among *h-continuation strategies*, fixing the part of his conjecture concerning types outside the coalition. This corresponds to requiring that any type in the coalition who plays a strategy which is ruled out by the restriction would prefer to make the restriction once h is reached, assuming that this does not affect what types outside the coalition play. Since it is required for every scenario that is consistent with h being reached, and for every type in the coalition, it is then unambiguously in the interest of the coalition to make any *h-based supported restriction* once h is actually reached. Making a *h-based supported restriction* then corresponds to the players foreseeing at the beginning of the game that this implicit agreement would be made if h is reached. Fixing the part of the conjecture that concerns types outside the coalition corresponds to the assumption that since the agreement is implicit, types outside the coalition cannot make their strategies contingent on whether the coalition made the restriction or not. Of course those players observe the previous action choices at any stage of the game and they update their beliefs at every stage. In particular they update their beliefs on whether types in the coalition made a particular agreement or not. The point is that *on top of what has been revealed by past actions* (and their strategies are already conditioned on that) there is no way that they can get information on whether the restriction has been made or not.

Note that the definition implies that if \underline{B} is an h -based supported restriction by Φ given \underline{A} and there is $\tau_i \in \Upsilon_i$ and $s_i \in \underline{A}_{\tau_i}$ such that $s_i \notin \underline{B}_{\tau_i}$ and $s_i \in BR_{\tau_i}(\theta_{-i})$ for some $\theta_{-i} \in \Theta_{-i}^c(\underline{A})$, then it has to be the case that \underline{A}_{-i} reaches h , otherwise (ii) cannot hold for type τ_i . Essentially, supported restrictions have to be based on histories that the participants consider to be possible to be reached.

Also note that in the payoff comparison at the definition of a supported restriction the expected payoff on the left hand side of the inequality (the one corresponding to a scenario when the restriction is violated) is required to be an updated expected payoff resulting from playing a best response strategy that reaches the given information set. On the right hand side though the payoff expectation is only required to result from playing a best response strategy *among replacement strategies* from the above information set. This is the formalization of the idea that a type in the coalition who plays a strategy outside the restriction (which is a best response against a conjecture) would always prefer to switch to making the agreement, and therefore switch to a new conjecture and play a best response to it among replacement strategies, at the history on which the agreement is based. Below we provide two examples to make clear some subtleties of this definition.




Figure 7

Consider the game of Figure 7. Player 3 has only one action choice in the first stage, so we omit it from the picture and from the description of stage-2 histories.

Note that playing $(A1, B1)$ after history (LL) (or excluding the play of $A2$ and $B2$ after history (LL)) is a supported restriction given S by players 1 and 2 (both players have just one type in this game). To check this, note that for both player 1 and 2 it holds that he only plays L in the first stage if he expects player 3 to play $C1$ after history (L, L) with at least probability $2/5$, because otherwise playing L in the first stage gives him a negative payoff, while playing R in the first stage guarantees a payoff of 0 for him. But if player 3 is expected to play $C1$ after history (L, L) with at least probability $2/5$, then $(A1, B1)$ gives a strictly higher payoff to both player 1 and 2 than any payoff they could obtain by playing $A2$ or $B2$. So in *all scenarios* in which history LL is reached, agreeing upon playing $(A1, B1)$ is mutually advantageous for players 1 and 2. This is true despite the fact that $A2$ and $B2$ are the unique best response continuation strategies continuation for players 1 and 2 (independently of what the other player does) after history (LL) if they believe that player 3 plays $C2$ with high enough probability. The point is that in those cases the history (LL) is never reached, therefore they do not matter for a coalitional agreement that concerns what to do after history (LL) occurred.




Figure 8

Consider now the game of Figure 8. Again we omit the trivial first stage move by player 3 in the description.

There is no supported restriction by any coalition in this game. In particular playing $(A1, B1)$ after history (LL) is not a supported restriction by players 1 and 2, because for example player 1 can play a strategy that plays L and then $B2$ if history (LL) occurred if he expects players 2 and 3 playing $B2$ and $C2$ with probability 1 after history (LL) . This gives an expected payoff of 2 to him. And in this case switching to an agreement to play $(A1, B1)$ after (LL) would give a strictly worse payoff expectation of 0, contradicting requirement (ii) in the definition of a supported restriction. Consider though modifying the previous requirement in the definition the following way:

- (i) \underline{B} is an h -based restriction by Φ given \underline{A}
- (ii) $\forall \theta_{-i} \in \Theta_{-i}^c(\underline{A}), \tau_i \in \Phi, s_i \in BR_{\tau_i}(\theta_{-i})$ it holds that $u_{\tau_i}^h(s_i, \theta_{-i}) < u_{\tau_i}^h(t_i, \psi_{-i})$
 $\forall t_i, \psi_{-i}$ such that $\psi_{-i} \in \Theta_{-i}^c(\underline{B}), \theta_{-\Phi}(h) = \psi_{-\Phi}(h), t_i \in A_{\tau_i}$ and t_i is a best response to ψ_{-i} .

In (ii') the expected payoff on the right hand side as well is required to come from playing a best response strategy (not just among h -replacement strategies) against a conjecture concentrated on \underline{B} , which reaches h . Modifying the definition this way makes playing $(A1, B1)$ after history (LL) a supported restriction by players 1 and 2. This is because if player 1 believes that player 2 plays L in the first stage with positive probability and then plays $B1$ after (LL) with probability 1, then he only plays L in the first stage if he believes that player 3 plays $C1$ with high enough probability. But if his beliefs are like that after (LL) , then no best response strategy can specify L in the first stage and then $A2$ after (LL) . Therefore there are no belief systems $\theta_{-1} \in \Theta_{-1}^c(\underline{A})$ and $\psi_{-1} \in \Theta_{-1}^c(\underline{B})$ that together satisfy the conditions in (ii'), making the latter trivially satisfied. Similarly for player 2. But clearly this just makes the agreement boot strapping! If players anticipate this agreement, then indeed in every case that they play L and therefore make it possible that history (LL) occurs, after (LL) they would agree upon the agreement. This does not make the agreement mutually advantageous though, even conditional on the given information set being reached. One can check for example that playing $(A2, B2)$ after history (LL) is similarly "supported" by players 1 and 2 according to the new definition. If players anticipate that agreement, then it is in their interest to make it after history (LL) .

Boot strapping agreements which are not unambiguously in the interest of participating types (players) do not satisfy the requirement in the definition of a supported restriction, which tests for whether it would be in the interest of players to make an agreement after a certain history, even if they didn't anticipate it beforehand.

Having defined supported restrictions for coalitions, given any product subset of strategies, we are ready to define the iterative procedure that defines the set of extensive form coalitionally rationalizable strategies. The procedure exactly corresponds to the one that defines the set of coalitionally rationalizable strategies in normal form. It starts out from the set of all possible strategies for types, \underline{S} . Then all strategies are discarded that are ruled out by some history based supported restriction by some coalition of types, given \underline{S} . In each subsequent step, all strategies are discarded that are ruled out by some history based supported restriction by some coalition of types, given the strategies that survived the previous step. The procedure corresponds to the assumption that at every step every coalition makes every history based restriction that is supported. Below we show that this assumption does not lead to contradictions and that at each step every type of every player has a nonempty set of strategies that is not discarded.

For every $\underline{A} = \prod_{i \in N} (\prod_{\tau_i \in \Upsilon_i} \underline{A}_{\tau_i})$ such that $\underline{A}_{\tau_i} \subset S_i$ and $\underline{A}_{\tau_i} \neq \emptyset \forall i \in N$ and $\tau_i \in \Upsilon_i$, let $F(\underline{A})$ be the set of supported restrictions given \underline{A} . Then let $\underline{A}^0 = \underline{S}$ and $\underline{A}^k = \bigcap_{B \in F(\underline{A}^{k-1})} B$. The decreasing sequence of sets $\underline{A}^0, \underline{A}^1, \dots$ represents iterated elimination of strategies that are never sequential coalitional best responses.

Definition: The set of extensive form coalitionally rationalizable strategies, denoted by \underline{A}^* is the set of strategies that survive iterated elimination of strategies that are never sequential coalitional best responses. $\underline{A}^* = \bigcap_{k \geq 0} \underline{A}^k$.

It is easy to show that in one-stage multi-stage games with complete information (each player has only one type) the set of extensive form coalitionally rationalizable strategies is exactly the same as the set of coalitionally rationalizable strategies of the normal form of the game. Since there is only one history (the null history) in these games, the concept of history based restrictions is equivalent in these games to the concept of restrictions. A belief process for any player is just a conjecture at the null history. Furthermore, a best response among replacement strategies from the null history to a certain conjecture is simply a best response strategy to the same conjecture. Therefore the definition of a supported restriction is essentially the same in the normal and the extensive form, and a restriction is supported in the extensive form game if and only if the corresponding restriction is supported in its normal form. This implies the equivalence of iterated elimination of strategies that are never sequential coalitional best responses in the original game with iterated elimination of strategies that are never coalitional best responses in its normal form. This establishes that extensive form coalitional rationalizability is indeed an extension of normal form coalitional rationalizability.

Now we show compatibility of history based supported restrictions. Claim 1 establishes that from any set of strategies that is reached by a nested sequence

of restrictions that is closed under rational behavior the intersection of all supported restrictions is nonempty, and adding it to the sequence gives another sequence that is closed under rational behavior. This is the key step in establishing nonemptiness of the set of extensive form coalitionally rationalizable strategies.

The proof establishes that for every type there is a belief process concentrated on the sequence of restrictions such that every best response strategy to this process has to be in every supported restriction by any coalition. Therefore the intersection of supported restrictions gives a nonempty set of strategies for every type. The above belief process can be created the following way. At the null history let the belief specified by the process be such that it yields at least as much expected payoff as any other belief that is concentrated on the sequence. At any subsequent history let the belief specified there be the same as the belief specified in its immediate predecessor history if the latter reaches the given history. Otherwise let it be a belief which gives the highest possible updated payoff at the history among those which reach the history and concentrated on the set in the sequence with the highest possible index. Informally this process is the “most optimistic” one that the type can have among the ones concentrated on the sequence. It starts out with the most optimistic belief and when there is a surprise, the belief is changed to the most optimistic one possible given the history that is reached. It is then not surprising that at no history it can be unambiguously in the interest of the type to agree upon giving up a best response strategy to this belief process.

All the proofs are in the appendix.

Claim 1: Let $(\underline{B}^1, \dots, \underline{B}^k) \in \mathcal{M}$ and $\underline{C} \equiv \bigcap_{\underline{B} \in F(\underline{B}^k)} \underline{B}$. Then $\underline{C} \neq \emptyset$ and $(\underline{B}^1, \dots, \underline{B}^k, \underline{C}) \in \mathcal{M}$.

Claim 2 establishes the central result of the paper, that the set of extensive form coalitionally rationalizable strategies is nonempty. It also states that the procedure of iterated elimination of strategies that are never sequential coalitional best responses stops after a finite number of sets and yields a sequence of nested restrictions that is closed under rational behavior.

Claim 2: $\underline{A}^* \neq \emptyset$ and $\exists K \geq 0$ such that $\underline{A}^* = \underline{A}^k \forall k \geq K$. Furthermore, $(\underline{A}^0, \dots, \underline{A}^K) \in \mathcal{M}$.

Claim 3 establishes that a restriction is supported by a one player coalition given some starting set if and only if the strategies that it rules out are never best responses to conjectures concentrated on the starting set.

Claim 3: Let $\underline{A} \subset \underline{S}$ such that $\underline{A} = \prod_{i \in N} \times_{\tau_i \in \Upsilon_i} A_{\tau_i} \neq \emptyset$. Let $i \in N$ and $\tau_i \in \Upsilon_i$. Then \underline{B} is a supported restriction by singleton coalition $\{\tau_i\}$ given \underline{A}

iff \underline{B} is a restriction by $\{\tau_i\}$ given \underline{A} and for every $s_i \in \underline{A}_{\tau_i}/\underline{B}_{\tau_i}$ it holds that there is no $\theta_{-i} \in \Theta_{-i}^c(\underline{A})$ such that $s_i \in BR_{\tau_i}(\theta_{-i})$.

Definition: The nested sequence $(\underline{B}^1, \dots, \underline{B}^k)$ is *coherent* if it is closed under rational behavior and for every $i \in N$, $\tau_i \in \Upsilon_i$ and $s_i \in \underline{B}_{\tau_i}^k$ it holds that $\exists \theta_{-i} \in \Theta_{-i}^c(\underline{B}^1, \dots, \underline{B}^k)$ such that $s_i \in BR_{\tau_i}(\theta_{-i})$.

A nested sequence is coherent if for every type of every player, the union of strategies that can be best responses against belief processes concentrated on the sequence is exactly the subset of strategies for the type that are included in each member of the sequence. It is a generalization of the concept of coherence in normal-form games (see Ambrus[01]). The next claim establishes that iterated elimination of strategies that are never sequential coalitional best responses yields a coherent sequence of restrictions.

Claim 4: Let $K \geq 0$ be such that $\underline{A}^K = \underline{A}^*$. Then $(\underline{A}^0, \dots, \underline{A}^K)$ is coherent.

The set of extensive form coalitionally rationalizable strategies is obtained via a specific iterative procedure, which makes every supported restriction by every coalition at each round. The next claim establishes that if players' beliefs concerning other players' types allocate positive probability to the actual realized type profile of the others, then the final outcomes that can be reached for the type profile is insensitive to the order of restrictions specified by the iterative procedure. There are two highlighted cases in which the former assumption holds. The first is when every player type allocates positive probability to every type profile of the other players, in which case the result holds for every type profile. A particular case of this is games of complete information (when every player has only one possible type). The second is when the beliefs of player types on other players' types are derived from a common prior probability distribution over the type space. In that case the above result holds for type profiles to which the common prior probability distribution allocates positive weight.

Claim 5: let $\underline{B}^0 = \underline{S}$. If there is no nontrivial supported restriction given \underline{B}^0 , then let $\underline{B}^1 = \underline{B}^0$. Otherwise let Ψ^0 be a nonempty collection of nontrivial restrictions given \underline{B}^0 and let $\underline{B}^1 = \bigcap_{\underline{B} \in \Psi^0} \underline{B}$. In a similar fashion once \underline{B}^k is defined for some $k \geq 1$, let $\underline{B}^{k+1} = \underline{B}^k$ if there is no nontrivial supported restriction given \underline{B}^k , otherwise let Ψ^k be a nonempty collection of nontrivial restrictions given \underline{B}^k and let $\underline{B}^{k+1} = \bigcap_{\underline{B} \in \Psi^k} \underline{B}$. Suppose $\tau \in \Upsilon$ is such that $\varphi_i(\tau_i)(\tau_{-i}) > 0 \forall i \in N$. Then there is L such that $O_\tau(\underline{B}^k) = O_\tau(\underline{A}^*)$, $\forall k \geq L$.

The next example demonstrates that the assumption that realized types' beliefs allocate positive probability to the actual realized type profile is necessary in Claim 5. The intuition is that the order of restrictions only has an affect on

possible action choices at histories that players never expect to be reached if they play the game with coalitionally rational opponents. But if a player type plays the game with other player types that he does not consider possible to be realized, then even if players are extensive form coalitionally rational, play can reach histories that she did not consider possible to be reached. And in that case the set of outcomes that can be reached depends on restrictions affecting histories that she considers impossible to be reached, which can depend on the order of restrictions.




Figure 9

In the game of Figure 1 player 1 has two types, I and II , while player 2 has only one possible type. For ease of exposition the picture only contains the public histories in the game (which are the same after both possible moves of nature). The payoff vectors at final histories are such that the first component denotes the payoff of type I of player 1, the second one denotes the payoff of type II of player 1 and the third denotes player 2's payoff. Assume that player 2 believes it with probability 1 that player 1's type is I . Then $\{Lr, Rl, Rr\} \times S_1 \times \{r \text{ after } L \text{ and } l \text{ after } R, r \text{ after } L \text{ and } r \text{ after } R\}$ is a $h = (L)$ based supported restriction for $1/I$ and 2 given \underline{S} . Furthermore, $S_1 \times S_1 \times \{r \text{ after } L \text{ and } r \text{ after } R, l \text{ after } L \text{ and } r \text{ after } R\}$ is a (null history based) supported restriction for the singleton coalition of 2 given \underline{S} . Therefore $\underline{A}^1 = \{Lr, Rl, Rr\} \times S_1 \times \{r \text{ after } L \text{ and } r \text{ after } R\}$. Then $\underline{A}^2 = \{Lr\} \times S_1 \times \{r \text{ after } L \text{ and } r \text{ after } R\}$ is a (null history based) supported restriction for the singleton coalition of $1/I$

given \underline{A}^1 . There is no more supported restriction in this game, so $\underline{A}^* = \underline{A}^2$. Note that if players play strategies in \underline{A}^* and the realized type of player 1 is II (which player 2 considers impossible), the the set of possible final histories is $\{(L, rl), (L, rr), (R, r)\}$. But now consider an alternative iterative procedure in which at the first round only the supported restriction $\underline{B}^1 = S_1 \times S_1 \times \{r \text{ after } L \text{ and } r \text{ after } R, l \text{ after } L \text{ and } r \text{ after } R\}$ by the singleton coalition of 2 is made. Then $\underline{B}^2 = \{Lr\} \times S_1 \times \{r \text{ after } L \text{ and } r \text{ after } R, l \text{ after } L \text{ and } r \text{ after } R\}$ is a supported restriction by the singleton coalition of 1/ I given \underline{B}^1 , and there is no supported restriction given \underline{B}^2 . But if players play strategies in \underline{A}^* and the realized type of player 1 is II , the the set of possible final histories is $\{(L, rl), (L, rr), (L, ll), (L, lr), (R, r)\}$, which is different from the final histories compatible with \underline{A}^* . The different order of restrictions changed the restrictions that affect only histories that players consider impossible to be reached at the end of the procedure. But if the realized types are such that one of these types allocates zero probability to the realized profile, then these histories can be reached by play and therefore it matters what restrictions were made at these histories at early stages of the iterative procedure.

Battigalli[97] defines correlated rationalizability in extensive form games as a natural modification of Pearce's extensive form rationalizability when one allows for correlated beliefs. Below this definition is stated using the terminology of this paper. Claim 6 establishes that if the realized types' beliefs allocate positive probability to the actual type profile of other players, then the final outcomes that are compatible with extensive form coalitional rationalizability is a subset of the outcomes that are compatible with extensive form correlated rationalizability⁹.

Let $\underline{R}^1 = \underline{S}$. For any $k \geq 1$ let $\underline{R}^k = \{\underline{s} \in \underline{R}^{k-1} \mid \forall i \in N, \tau_i \in \Upsilon_i \exists \theta_{-i} \in \Theta_{-i}^c(\underline{R}^{k-1}) \text{ such that } s_{\tau_i} \in BR_{\tau_i}(\theta_{-i})\}$. The decreasing set of strategies $(\underline{R}^0, \underline{R}^1, \dots)$ represents iterated deletion of strategies that are never weak sequential best responses to correlated conjectures.

Definition: the set of extensive form rationalizable strategies is $\underline{R}^* \equiv \bigcap_{k=0,1,2,\dots} \underline{R}^k$.

For more on extensive form rationalizability, see Pearce[84] and Battigalli[97].

Claim 6: Suppose $\tau \in \Upsilon$ is such that $\varphi_i(\tau_i)(\tau_{-i}) > 0 \forall i \in N$. Then $O_\tau(\underline{A}^*) \subset O_\tau(\underline{R}^*)$.

⁹similarly, if extensive form coalitional rationalizability is defined on independent conjectures, then the outcomes compatible with it is a subset of the outcomes compatible with extensive form rationalizability (defined on independent conjectures).

4 Relationship to extensive form equilibrium concepts

4.1 Perfect coalition-proof and renegotiation-proof equilibria

Ambrus[01] establishes that in normal form games there is no inclusion relationship between the set of profiles that can be part of some coalition-proof Nash equilibrium and the set of coalitionally rationalizable strategies. It is easy to extend this result to extensive form rationalizability, and perfect coalition-proof Nash equilibrium (see Bernheim, Peleg and Whinston[87]). Since, as we showed in section 3, extensive form coalitional rationalizability is a generalization of normal form coalitional rationalizability, and perfect coalition-proof Nash equilibrium is a generalization of coalition-proof Nash equilibrium, the examples provided in the above paper also establish that there is no inclusion relationship between the set of profiles that can be part of some perfect coalition-proof Nash equilibrium and the set of extensive form coalitionally rationalizable strategies.

But there are differences between perfect coalition-proof Nash equilibrium and extensive form coalitional rationalizability that are caused by the differences of imposing sequential (coalitional) rationality in an equilibrium framework and in a non-equilibrium framework. Ambrus[01] shows that if every restriction of a normal form game has a coalition-proof Nash equilibrium, then all coalition-proof Nash equilibria of the game are contained in the set of coalitionally rationalizable strategies. In particular every coalition-proof Nash equilibrium of every two-player game is contained in the set of coalitionally rationalizable strategies. Below we provide an example of a two-player extensive form game in which a perfect coalition-proof Nash equilibrium is not contained in the set of extensive form rationalizable, and therefore in the set of extensive form coalitionally rationalizable strategies. The reason is that perfect coalition-proof Nash equilibrium does not entail the type of forward induction reasoning implied by extensive form rationalizability and extensive form coalitional rationalizability.

One difficulty involved in comparing perfect coalition proof Nash equilibria and the set of coalitionally rationalizable strategies is that sequential considerations in perfect coalition-proof Nash equilibrium, as defined in Bernheim, Peleg and Whinston[87], are subgame based. The given candidate profile has to be a (perfect) coalition-proof Nash equilibrium in every proper subgame of the original game. This means that perfect coalition-proof Nash equilibrium does not refine coalition-proof Nash equilibrium in multi-stage games in which at least one player has more than one type. Although it is possible to provide a modification of the definition of coalition-proof Nash equilibrium which is effective

in incomplete information games,¹⁰ we do not pursue that direction here and analyze a game with complete information below.

	B1	B2
A1	100,10	9,0
A2	0,9	10,100
A3	1,0	1,0

Figure 10

Consider the twice repeated version of the game of Figure 10. The claim below establishes that extensive form rationalizability implies that players have to play $(A1, B1)$ in the second stage if the realized history in the first stage is $(A2, B2)$. Since by Claim 6 the final outcomes consistent with extensive form coalitional rationalizability is a subset of the final outcomes consistent with extensive form rationalizability, this establishes that the strategy of choosing $A2$ in the first period and then $A2$ independently of the first period history is not coalitionally rationalizable for player 1. Similarly, the strategy of choosing $B2$ in the first period and then $B2$ independently of the first period history is not coalitionally rationalizable for player 2. But note that the above profiles constitute a perfect coalition-proof Nash equilibrium. In every second period subgame $(A2, B2)$ is played, which is a coalition-proof Nash equilibrium. And there is no Nash equilibrium which Pareto dominates the above equilibrium, since it yields the highest possible payoff for player 2, and therefore the above profile is a perfect coalition-proof Nash equilibrium.¹¹

Denote the above game by \tilde{G} .

Claim 7: There is no extensive form rationalizable strategy of player 1 in \tilde{G} which specifies choosing $A2$ with positive probability in the first stage and then choosing $A2$ with positive probability after a first stage history $(A2, B2)$. Furthermore, There is no extensive form rationalizable strategy of player 2 in \tilde{G} which specifies choosing $B2$ with positive probability in the first stage and then choosing $B2$ with positive probability after a first stage history $(A2, B2)$.

¹⁰Kahn and Mookherjee[95] offer a definition of coalition-proof Nash equilibrium in a class of games with private information. But the games they consider differ from standard Bayesian games in that all types of all players are actually present in the game and their action choices are payoff relevant and different types of the same player are only connected by being indistinguishable for the other players. Their construction also assumes explicit communication for coalitional deviations, while we concentrate on the case when no communication is possible among players during the game.

¹¹the above profile is also strong perfect Nash equilibrium, according to the definition provided in Rubinstein[80] applied to finitely repeated games, showing that not all strong perfect Nash equilibria can be extensive form coalitionally rationalizable.

The intuition behind Claim 7 is the following. It can be shown that playing $A2$ or $A3$ in the first period can only be optimal for player 1 if he expects player 2 to play $B2$ with positive probability in that period. This implies that player 1 cannot be surprised after histories $(A2, B2)$ and $(A3, B2)$ and therefore playing $A2$ or $A3$ can be an unambiguous signal of strategic intent of what he plans to play in the second period. And it is indeed the case, because playing $A3$ in the first period can only be rational for him if he plans to play $A1$ after history $(A3, B2)$. Therefore the only rationalizable second period outcome after a first stage history $(A3, B2)$ is $(A1, B1)$. But it is possible to show that then it can never be optimal for player 1 to play a strategy that plays $A2$ in the first period and then not $A1$ in the second period if the realized history was $(A2, B2)$. Therefore the only rationalizable second period outcome after a first stage history $(A2, B2)$ is $(A1, B1)$.

Forward induction implied by extensive form rationalizability or related concepts of iterated domination is analyzed extensively in the literature on games where a player has an outside option at the beginning of the game and in “money-burning” games (see for example Ben-Porath and Dekel[92] and Shimoji[02]). Furthermore, equilibrium refinements that incorporate forward induction logic were proposed and analyzed in games where players can potentially move simultaneously in every stage of the game (see Kohlberg and Mertens[86], VanDamme[89] and Osborne[90]). In games in which players move simultaneously at every stage of the game, in particular in finitely repeated games, extensive form rationalizability is less effective in restricting the set of strategies using forward induction considerations, because if a player can be surprised by the other players’ actions, then his past actions might not reflect strategic intent in future periods. But as the game of Figure 10 demonstrates, extensive form rationalizability can imply forward induction type restrictions even in these games.

Since in two-player finitely repeated games perfect coalition-proof Nash equilibrium coincides with renegotiation-proof Nash equilibrium (see Benoit and Krishna[93], Bernheim and Ray[89] and Bernheim, Peleg and Whinston[87]), the above example also demonstrates that even in two-player games not all renegotiation-proof Nash equilibria have to be included in the set of extensive form rationalizable, and therefore in the set of extensive form coalitionally rationalizable strategies. This observation is less relevant if renegotiation-proofness is considered as an extra requirement in contexts where players indeed get together and explicitly negotiate over continuation strategies. If that is the case, forward induction considerations implied by extensive form rationalizability and extensive form coalitional rationalizability might lose their force, since a player can change his mind during the negotiation phase and therefore his action choices earlier in the game might not reflect strategic intent in the remaining game. This is the main reason we consider extensive form rationalizability to be a relevant solution concept in multi-stage games if players cannot communicate during the game. But the above example challenges the view that the concept

of renegotiation-proofness is valid even if players do not explicitly negotiate over strategies after each stage of the game¹².

4.2 Perfect equilibrium concepts

It is easy to find examples that show that there is no containment relationships between the set of extensive form coalitionally rationalizable strategies on one

hand and the set of perfect Bayesian equilibria, sequential equilibria and trembling hand perfect equilibria¹³. Furthermore, we show below that it is possible that no trembling hand perfect equilibrium of a game is included in the set of extensive form coalitionally rationalizable strategies. However, we show that every finite multi-stage game has at least one sequential equilibrium which is fully contained in the set of extensive form coalitionally rationalizable strategies.




Figure 11

¹²this view is suggested for example in Benoit and Krishna[93] by their statement “We emphasize that the assumption of explicit communication is made for the purposes of motivation only...”.

¹³given that extensive form coalitional rationalizability is an extension of normal form coalitional rationalizability (as established in the previous section) the examples provided in Ambrus[01] that show that in normal form games not every coalitionally rationalizable outcome profile has to be part of some Nash equilibrium and that not every Nash equilibrium is contained in the set of coalitionally rationalizable strategies suffice here, since in simultaneous-move games every Nash equilibrium is trivially sequential and trembling-hand perfect equilibrium.

In the game of Figure 11 the set of extensive form coalitionally rationalizable strategies is the same as the set of extensive form rationalizable strategies: $\{OL, OR, IR\} \times \{m, r\}$. In the first round of iterated deletion of strategies that are never sequential (coalitional) best responses, strategy IR of player 1 is eliminated, and in the second round strategy l of player 2 is deleted, resulting in the above set. This corresponds to a standard forward induction reasoning. But the unique trembling hand perfect equilibrium of the game is (OL, l) , in which player 2 plays a strategy that is not extensive form rationalizable. It is because against any strictly mixed strategy of player 1, r can never be a best response to player 2 (who has only one information set and therefore one agent in the agent normal form of the game). But then in every trembling hand perfect equilibrium player 1 has to choose L at her second information set (decision node), which in turn implies that she has to play O at her first information set and player 2 has to play l .

The example shows the difference between how extensive form coalitional rationalizability (and extensive form rationalizability) and trembling hand perfect equilibrium gives interpretation to players to a history being reached. According to extensive form coalitional rationalizability, player 1 choosing I at the beginning of the game should be interpreted as a signal that she plans to play R at her second decision node and that justifies player 2 playing r . According to trembling hand perfection there is some (infinitesimal) chance that this move is a result of a mistake, and furthermore there is some chance of player 1 choosing either of the two actions at her second decision node. Therefore in the context of trembling hand perfection player 2 should never play r . The forward induction reasoning implied by extensive form coalitional rationalizability assumes that players try to attach a rational explanation to every observed history. Trembling hand perfection invalidates this type of reasoning and as shown above can lead to predictions that are not compatible with extensive form coalitional rationalizability.

Sequential equilibrium (see Kreps and Wilson[82]) is a refinement of Nash equilibrium in extensive form games that is defined by putting restrictions on players' conjectures and is therefore much closer to the framework provided in this paper. In that sense it is not surprising that as opposed to trembling hand perfect equilibrium, every multi-stage game has a sequential equilibrium that is fully contained in the set of extensive form coalitional rationalizable strategies. This also implies that every multi-stage game has a perfect Bayesian equilibrium that is fully contained in the set of extensive form coalitional rationalizable strategies.

To be compatible with traditional equilibrium analysis, in the following claim we make the assumption that players' conjectures on types come from a common prior probability distribution on types¹⁴.

¹⁴this is not essential. The claim holds for subjective sequential equilibrium in multi stage games with incomplete information without the common prior assumption.

Let $\Sigma = \times_{i \in N} \Sigma_i$ be the set of mixed strategies.

Claim 8: there is $\sigma \in \Sigma$ such that σ is a sequential equilibrium of G and $\text{supp} \sigma \subset \underline{A}^*$.

5 Special classes of multi-stage games

5.1 Complete information, perfect information

This subsection considers the subset of multi-stage games in which information is complete. The multi-stage game G is of complete information if every player has only one type: $\|\Upsilon_i\| = 1 \forall i \in N$. To simplify notation, we associate the only element of Υ_i with i , $\forall i \in N$. Therefore any coalition of types $\Phi = \{\tau_j \mid j \in J\}$ will be associated with the coalition of players J . Furthermore, since $\underline{S} = S$, only the latter notation is used. Similarly, a typical subset of $\underline{S} = S$ is denoted by A instead of \underline{A} . Since there is a unique profile of types in the game, for every $A \subset S$ we just use $O(A)$ to denote the set of final histories reached by strategy profiles in A , instead of $O_\tau(A)$.

The definition of a history-based supported restriction is then simplified as follows.

Let $A = \times_{i \in N} A_i$ be such that $A \subset S$ and $A \neq \emptyset$. Let $J \subset N$. Let $h \in H/H^Z$.

Definition: B is an h -based supported restriction by J given A if

- (i) B is an h -based restriction by J given A
- (ii) $\forall j \in J$, $s_j \in A_j/B_j$ and $\theta_{-j} \in \Theta_{-j}^c(A)$ such that $s_j \in BR_j(\theta_{-j})$ it holds that $u_j^h(s_j, \theta_{-j}) < u_j^h(t_j, \psi_{-j})$
 $\forall t_j, \psi_{-j}$ such that $\psi_{-j} \in \Theta_{-j}^c(B)$, $\theta_{-j}(h) = \psi_{-j}(h)$, $t_j \in A_j$ and t_j is a best

response to ψ_{-j} among h' -replacement strategies of s_j for every $h' \in h \cup \text{succ}(h)$ that is reached by t_j .

A special case of complete information is perfect information. A multi-stage game is of perfect information if besides every player having only one possible type, at every history there is at most one player who has more than one action choice at the history.

Definition: G is of perfect information if $\|\Upsilon_i\| = 1 \forall i \in N$ and at every $h \in H/H^Z$ if $\|A_i(h)\| \geq 2$ and $\|A_j(h)\| \geq 2$ for $i, j \in N$, then $i = j$.

Claim 9 below establishes that in games of perfect information extensive form coalitional rationalizability is outcome-equivalent to extensive form rationalizability. In games of perfect information without relevant ties extensive form rationalizability is outcome equivalent to the unique backward induction solution (see for example Battigalli[97] for the definition of games without relevant ties and the proof of the above claim) and then Claim 6 immediately implies outcome equivalence of extensive form rationalizability and extensive form coalitional rationalizability. But Claim 9 establishes the equivalence result for every game of perfect information. This means that in games of perfect information there is no room for non-singleton coalitions to make credible agreements. The intuition is that extensive form rationalizability only does not pin down a unique outcome in a perfect information game if at some relevant information set the player who moves there is indifferent between more than one action choices. So any coalitional agreement would involve breaking ties. This might be strictly advantageous for some players. But at the last information set affected by the restriction the player who moves there does not have a strict incentive not to play a strategy outside the restriction, which undermines the credibility of the restriction.

Claim 9: If G is a game of perfect information, then $O(A^*) = O(R^*)$.

5.2 Normal form Bayesian games

A multi-stage game is called a normal form Bayesian game if the only nonterminal history is the null history. In these games every restriction is null-history based. Belief processes are trivial, $\Theta_{-i} = \Theta_{-i}^c = \Omega_{-i} \forall i \in N$. This simplifies the definition of a supported restriction the following way.

Let $\underline{A} = \prod_{i \in N} (\prod_{\tau_i \in \Upsilon_i} \underline{A}_{\tau_i})$ be such that $\underline{A}_{\tau_i} \subset S_i$ and $\underline{A}_{\tau_i} \neq \emptyset \forall i \in N$ and $\tau_i \in \Upsilon_i$. Let $J \subset N$ and $\Phi_j \subset \Upsilon_j \forall j \in J$. Let $\Phi = \bigcup_{j \in J} \Phi_j$.

Definition: \underline{B} is a supported restriction by Φ given \underline{A} if

- (i) $\underline{B} = \prod_{i \in N} (\prod_{\tau_i \in \Upsilon_i} \underline{B}_{\tau_i}) \subset \underline{A}$, $\underline{B} \neq \emptyset$ and $\underline{B}_{\tau_i} = \underline{A}_{\tau_i} \forall i \in N$ and $\tau_i \notin \Phi$.
- (ii) $\forall \tau_i \in \Phi$, $s_i \in \underline{A}_{\tau_i} / \underline{B}_{\tau_i}$ and $\omega_{-i} \in \Omega_{-i}(\underline{A})$ such that $s_i \in BR_{\tau_i}(\omega_{-i})$ it holds that $u_i(\tau_i, s_i, \omega_{-i}) < u_i(\tau_i, t_i, \omega'_{-i})$
 $\forall t_i, \omega'_{-i}$ such that $\omega'_{-i} \in \Omega_{-i}(\underline{B})$, $\omega_{-\Phi} = \omega'_{-\Phi}$, $t_i \in \underline{A}_{\tau_i}$ and $t_i \in BR_{\tau_i}(\omega'_{-i})$.

This definition is similar to the definition of supported restriction in normal form games, with types replacing players. Claim 10 below makes this connection formal and provides insight on what coalitional agreements are supported in an incomplete information environment. It establishes that the set of extensive

form coalitionally rationalizable strategies of a Bayesian game exactly corresponds to the set of normal form rationalizable strategies of the game obtained by letting types of players in the original game be players in the normal form game and defining payoffs appropriately.¹⁵

Let $\widehat{G} = (\widehat{N}, \widehat{S}, \widehat{u})$ be the normal form game obtained from G as follows. The set of players is $\widehat{N} = \bigcup_{i \in N} \Upsilon_i$, the set of strategies for every $\tau_i \in \widehat{N}$ is $\widehat{S}_{\tau_i} = S_i$ (therefore $\widehat{S} = \underline{S}$), and the payoff function of every $\tau_i \in \widehat{N}$ and $\underline{s} \in \prod_{\tau_i \in \widehat{N}} S_{\tau_i}$ is
$$\widehat{u}(\underline{s}) = \sum_{\tau'_{-i} \in \Upsilon_{-i}} u_i(\tau_i, s_{\tau'_1}, \dots, s_{\tau'_{i-1}}, s_{\tau_i}, s_{\tau'_{i+1}}, \dots, s_{\tau'_N}) \cdot \varphi^{\tau_i}(\tau'_{-i}).$$

The payoff of player τ_i in the above definition is independent of the strategies played by players corresponding to other types of the same player i in the original game G . Otherwise the payoff of player τ_i is a weighted average of payoffs from G , corresponding to possible type profiles of the other players in G . The weights used are the probabilities that the posterior belief of type τ_i allocates to possible type profiles of the other players in G . This corresponds to a scenario in which every realized player type knows that other player types of the same player are not realized, and his beliefs on what types of other players are realized are given by his posterior belief concerning types.

Let \underline{A}'^* denote the set of coalitionally rationalizable strategies of G' .

Claim 10: Let $\underline{A} \subset \underline{S}$ such that $\underline{A} = \prod_{i \in N} \left(\prod_{\tau_i \in \Upsilon_i} \underline{A}_{\tau_i} \right) \neq \emptyset$. Let $\underline{B} \subset \underline{A}$ such that $\underline{B} = \prod_{i \in N} \left(\prod_{\tau_i \in \Upsilon_i} \underline{B}_{\tau_i} \right) \neq \emptyset$. Then \underline{B} is a supported restriction by $J \subset N$ given \underline{A} in Bayesian game G iff it is a supported restriction by J given \underline{A} in normal form game G' . Furthermore, $\underline{A}'^* = \underline{A}^*$.

6 Ex ante coalitional agreements and sequential rationality

Extensive form coalitional rationalizability takes the view that groups of players can make supported restrictions conditional on information sets being reached and conditional on players' types. This means that new coalitional agreements can be made as the game progresses, after nature selected players' types and after any observed history of play. We consider this construction to be the natural extension of sequential rationality to the context of coalitional agreements.

¹⁵Imposing the additional restriction that different types of the same player have to have the same belief processes over other players' strategies would of course break down this equivalence.

But as the motivating examples of Figure 5 and Figure 6 in section 2 show, coalitional agreements in this setting might not be advantageous for the players involved from an ex ante point of view. Therefore if players can commit to making coalitional agreements only at the beginning of the game, they might choose to do so.

This section provides an alternative construction, in which coalitional agreements can only be made at the beginning of the game and are evaluated from an ex ante perspective. The latter corresponds to the assumption that players can only make implicit coalitional agreements that are advantageous for the players involved *before their types are realized*. The assumption of sequential individual rationality is maintained though, and it is allowed to interact with ex ante coalitional agreements. The game of Figure 12 below demonstrates that this interaction is nontrivial, in the sense that iterative deletion of strategies that are not sequential best responses might make ex ante coalitional restrictions supported and then that restriction might make new strategies never sequential best responses.

To be able to make ex ante payoff comparisons, in this section it is assumed that there is a common prior probability distribution over the set of types, $\varphi(\tau)$.

For every $i \in I$ and $\tau_i \in \Upsilon_i$ let $\varphi^{\tau_i} = \sum_{\substack{v \in \times_{j \in N} v_j: v_i = \tau_i}} \varphi(v)$, the marginal probability of type τ_i of player i being realized. Let h_0 denote the null history. Let $\underline{A} = \times_{i \in N} (\times_{\tau_i \in \Upsilon_i} \underline{A}_{\tau_i})$ be such that $\underline{A}_{\tau_i} \subset \underline{S}_i$ and $\underline{A}_{\tau_i} \neq \emptyset \forall i \in N$ and $\tau_i \in \Upsilon_i$. Let $J \subset N$.

Definition: \underline{B} is an ex ante supported restriction by J given \underline{A} if

- (i) $\underline{B} = \times_{i \in N} (\times_{\tau_i \in \Upsilon_i} \underline{B}_{\tau_i}) \subset \underline{A}$, $\underline{B} \neq \emptyset$ and $\underline{B}_i = \underline{A}_i \forall i \notin J$.
- (ii) $\forall j \in J$, $\theta_{-j} \in \Theta_{-j}^c(\underline{A})$ and $\underline{s} : \Upsilon_j \rightarrow S_j$ for which $\underline{s}_{\tau_j} \in \underline{A}_{\tau_j}$ and $\underline{s}_{\tau_j} \in BR_{\tau_j}(\theta_{-j}) \forall \tau_j \in \Upsilon_j$ and there is $\tau_j \in \Upsilon_j$ such that $\underline{s}_{\tau_j} \notin B_{\tau_j}$ it holds that $\sum_{\tau_j \in \Upsilon_j} \varphi^{\tau_j} \cdot u_j(\tau_j, \underline{s}_{\tau_j}, \theta_{-j}) < \sum_{\tau_j \in \Upsilon_j} \varphi^{\tau_j} \cdot u_j(\tau_j, \underline{t}_{\tau_j}, \theta'_{-j})$
 $\forall \theta'_{-j} \in \Theta_{-j}^c(\underline{B})$ and $\underline{t} : \Upsilon_j \rightarrow S_j$ for which $\underline{t}_{\tau_j} \in \underline{A}_{\tau_j}$ and $\underline{t}_{\tau_j} \in BR_{\tau_j}(\theta'_{-j})$
 $\forall \tau_j \in \Upsilon_j$, and $\theta'_{-j}(h_0) = \theta_{-j}(h_0)$.

A restriction is supported by a coalition of players if every player in the coalition has a strictly higher ex ante expected payoff if he expects the restriction to be made (if his belief process is concentrated on the restriction) than if at least one of his types plays outside the restriction (if he has a belief process against which at least one of his types has a best response outside the restriction), holding null-history conjectures concerning players outside the coalition fixed. The definition (just like the definition of a history based supported restriction

in section 3) implies that an ex ante supported restriction is self-enforcing. For every type of every player it holds that if he believes that the other players play according to the agreement, then a strategy that is ruled out by the agreement can never be a best response for him.

For every $\underline{A} = \prod_{i \in N} (\prod_{\tau_i \in \Upsilon_i} \underline{A}_{\tau_i})$ be such that $\underline{A}_{\tau_i} \subset S_i$ and $\underline{A}_{\tau_i} \neq \emptyset \forall i \in N$ and $\tau_i \in \Upsilon_i$ let $F^e(\underline{A})$ be the set of ex ante supported restrictions given \underline{A} , and let $R(\underline{A})$ be such that $R(\underline{A}) = \prod_{i \in N} (\prod_{\tau_i \in \Upsilon_i} R_{\tau_i}(\underline{A}))$ where $R_{\tau_i}(\underline{A}) = \bigcup_{\theta_{-j} \in \Theta_{-j}^c(\underline{A})} (BR_{\tau_i}(\theta_{-j}) \cap A_{\tau_i})$. For every $i \in N$ and $\tau_i \in \Upsilon_i$ $R_{\tau_i}(\underline{A})$ is the set of strategies in A_{τ_i} which can be best responses to some belief process concentrated on \underline{A} .

Let $\underline{E}^0 = \underline{S}$. For $k \geq 1$ define \underline{E}^k iteratively as $\underline{E}^k = \bigcap_{B \in F^e(\underline{E}^{k-1})} B \cap R(\underline{A})$.

The set of strategies \underline{E}^k is obtained from \underline{E}^{k-1} by making all ex ante supported restrictions given \underline{E}^{k-1} and deleting all strategies in \underline{E}^{k-1} that are never (weak sequential) best responses against belief processes concentrated on \underline{E}^{k-1} .

Definition: the set of ex ante coalitionally rationalizable strategies is $\underline{E}^* = \bigcap_{k \geq 0} \underline{E}^k$.

The set of ex ante coalitionally rationalizable strategies have similar properties as the set of extensive form coalitionally rationalizable strategies. In particular $\underline{E}^* \neq \emptyset$, there is $L \geq 0$ such that $\underline{E}^k = \underline{E}^L = \underline{E}^* \forall k \geq L$, the sequence of restrictions $(\underline{E}^0, \dots, \underline{E}^L)$ is coherent, and $O(\underline{E}^*) \subset O(\underline{E}^*)$. The proofs of these claims are similar to the proofs of corresponding claims concerning \underline{A}^* in section 2 and therefore omitted.




Figure 12

The game of Figure 12 shows that there can be nontrivial interaction between ex ante supported restrictions and elimination of never sequential best response strategies.

In this game no coalition has an ex ante supported restriction given \underline{S} . But any strategy which specifies action r for player 3 after history (r, r) is never a sequential best response for her. Eliminating those strategies gives $\underline{E}^1 = S_1 \times S_2 \times \{lL, rL\}$. Every strategy for every player in \underline{E}^1 is a best response to some belief process concentrated on \underline{E}^1 . But $\{lL, mL, mR\} \times \{lL, mL, mR\} \times \{lL, rL\}$ is an ex ante supported restriction for the coalition of player 1 and player 2 given \underline{E}^1 . There is no other ex ante supported restriction given \underline{E}^1 , therefore $\underline{E}^2 = \{lL, mL, mR\} \times \{lL, mL, mR\} \times \{lL, rL\}$. There is no ex ante supported restriction given \underline{E}^2 . But rL is not a sequential best response for player 3 to any belief process concentrated on \underline{E}^2 . Every other strategy for every player in \underline{E}^2 can be a best response to some belief process concentrated on \underline{E}^2 and therefore $\underline{E}^3 = \{lL, mL, mR\} \times \{lL, mL, mR\} \times \{lL\}$. There is no ex ante supported restriction given \underline{E}^3 and every strategy for every player in \underline{E}^3 can be a best response to some belief process concentrated on \underline{E}^3 and therefore $\underline{E}^* = \underline{E}^3$.

Note that in the above game iterated elimination of never sequential best response strategies and ex ante supported restrictions are intertwined in a non-trivial manner. A procedure of first taking the extensive form rationalizable strategies of the game and then iteratively making all possible ex ante supported restrictions does not eliminate all outcomes that are eliminated above. Similarly, a procedure of first iteratively making all possible ex ante supported restrictions and then iteratively eliminating strategies that are never sequential best responses fails to eliminate all outcomes that are not compatible with ex ante coalitional rationalizability.

7 Conceptual issues in general extensive form games

In multi-stage games with observable actions players can condition a coalitional agreement on a history being reached because it is public knowledge if the history was reached and so are players' possible choices in the remainder of the game (the continuation strategies from that history). In general extensive form games it is less obvious what should be the right conditioning events for coalitional agreements and how (from the point of view of which information set) players should evaluate these agreements. We do not attempt to provide a satisfactory answer to these questions in this paper, just provide two examples below to show some of the difficulties associated with sequential coalitional reasoning in general extensive form games.




Figure 13

Consider the restriction $\{L\} \times \{I, O\} \times \{l\}$ by $\{1, 2\}$ in the game of Figure

13. In words, an agreement between players 1 and 3 to play L and l respectively. Is this a mutually advantageous agreement? It certainly is for player 1, since the agreement guarantees him a payoff of at least 2, while if he plays R , he can never get a payoff higher than 1. As far as player 3 is concerned though, the answer depends on whether he evaluates this agreement from the point of view of the beginning of the game or from the point of view of his information set. If the agreement is evaluated from the point of view of the beginning of the game, then the first problem with the agreement is that player 3's information set might not be reached and therefore it might be immaterial whether he plans to play l or r . Second, if player 2 plays O with high probability, then he gets a higher payoff if player 1 plays R and he plays r than if they play according to the agreement. So the agreement does not seem to be unambiguously advantageous for player 3. But now consider the following argument. Player 3 should evaluate the agreement $\{L\} \times \{I, O\} \times \{l\}$ from the point of view of his information set, because whether he would agree upon making the agreement and play l is only relevant if his information set is reached. But then note that conditional on his information set being reached player 3's highest payoff is when player 1 played L and he plays l . If his information set is reached, player 1 would agree to make the above agreement and play l , because it becomes irrelevant for him what happens in scenarios when his information set is not reached.




Figure 14

Compatibility of coalitional agreements becomes a tricky issue as well in general extensive form games, as the game of Figure 14 demonstrates. The

restriction $\{R\} \times S_2 \times \{l\}$ (agreeing upon playing R and l) given S seems to be a mutually advantageous one for players 1 and 3, since for every fixed mixed strategy of player 2 this agreement guarantees the highest possible expected payoff for both players 1 and 3. Similarly, the agreement $\{L, R\} \times \{I \text{ after } L \text{ and } O \text{ after } R\} \times \{r\}$ (player 2 agreeing upon playing I after L and O after R , and player 3 agreeing upon playing l) seems to be a mutually advantageous one for players 2 and 3, because for every fixed strategy of player 1 it guarantees the highest possible expected payoff for both players 2 and 3. But note that if both of these agreements are made, then player 3 made an agreement with player 1 not to play l , and he made an agreement with player 2 not to play r ! Therefore the two agreements seem to be incompatible. On the other hand note that if these agreements are made then player 3's information set is never reached. Therefore he can promise to player 1 not to play r and to player 2 not to play l , because he is guaranteed not to get into a position in which he had to break one of the promises. This suggests to consider coalitional agreements in general extensive form games with the feature that players agree upon restricting their set of strategies as long as play does not reach an information set which is incompatible with the restrictions (as long as the player did not receive evidence that some other player violated an agreement).

Because of the above described complications arising from allowing players in a coalition to evaluate the same agreement from different information sets, even if allowing for that seems natural in general extensive form games, one might want to insist on players evaluating the agreement from the same point of view. In the above two games the natural point would be the beginning of the game. In general the right common conditioning event for evaluating the agreement would have to be an event which, informally speaking, is compatible with all players knowing that none of the information sets "affected by the agreement" have been reached yet. This claim is of course just informal without defining formally the information sets affected by a restriction. Formalizing this requirement and investigating whether there is a maximal conditioning event like that is left to future investigation. This question is particularly involved in games without an explicit time structure, where the same information sets can be reached both before and after the ones that are "affected by the agreement", unless the definition of information sets affected by a restriction automatically implies that once an information set is affected by a restriction then all information sets that can possibly succeed it are affected by the restriction too.

8 Related literature

Section 4 provides references to papers proposing equilibrium refinements in extensive form games that are related to the concept of coalitional agreements. The same section provides references to the literature on forward induction

implied by inferences on strategic intent from past actions of a player. This section gives an overview of other segments of the literature that are connected to extensive form coalitional rationalizability.

Battigalli and Siniscalchi[01] proposes a class of refinements of extensive form rationalizability in games of incomplete information with observable actions, by explicitly imposing additional assumptions on first order beliefs. They collectively call the resulting concepts strong Δ -rationalizability. By Claim 4 of Section 3 above the sequence of restrictions $(\underline{A}^0, \dots, \underline{A}^*)$ is coherent and therefore extensive form coalitional rationalizability can be regarded as a result of extra restrictions on first order beliefs - players' conjectures are required to be concentrated on the nested sequence $(\underline{A}^0, \dots, \underline{A}^*)$. Battigalli and Siniscalchi consider more general belief processes (conditional probability systems in their terminology) than those we allow for. Belief processes in their framework are defined over both types and strategies of other players. This implies that after being surprised, a player (type) is allowed to change his belief on the original distribution over other players' types too, not only on their strategies. Furthermore, correlation is allowed between the beliefs concerning types and strategies. But it is easy to show that our framework results from putting extra assumptions on first order beliefs in their general framework, and therefore extensive form coalitional rationalizability is a strong Δ -rationalizability concept.

Extensive form coalitional rationalizability is a theory of implicit agreements by players along the course of play. A complementary research agenda is characterizing reasonable outcomes in extensive form games if players have explicit communication possibilities. Myerson[89] investigates Bayesian games in which at most one player has more than one possible type and only this player can send a message to the other players before playing the game. Because of the special features of the model, the solution concept proposed does not treat players symmetrically. In particular groups of players that do not involve the player who can send messages do not have the opportunity to make agreements with each other. Myerson[86] combines the concept of communication equilibrium (see for example Forges[86]) with the requirement of sequential rationality in multi-stage games with communication mechanisms. His paper is primarily concerned with how communication can expand the possibilities of players, subject to sequential rationality constraints, and not with the issue of what agreements will actually be made by players. A major difference between our framework and that of the above papers' is that communication mechanisms introduce the possibility of correlated play, while in this paper we allow for correlated conjectures, but not correlated play.

Extensive form coalitional rationalizability implies that the implicit agreements that different types of players make are self-enforcing in the sense that if a player type assumes that other types play according to the agreement, then it is strictly in his interest to play according to the agreement. Holmstrom and Myerson[83] provides analysis of incentive compatibility in the context of

incomplete information in a cooperative environment. The notion of the incentive compatible core (see for example Vohra[99] and Dutta and Vohra[01] and Forges, Minelli and Vohra[02]) combines this analysis with coalitional contracts through the concept of the core. The cooperative framework enables players to make binding agreements with the constraint that private information is not verifiable. Our paper on the other hand assumes that players cannot make any kind of binding agreements before or during the play of the game. The issue of incentive compatibility in a noncooperative incomplete information contracting game is taken up by Maskin and Tirole[90] and Maskin and Tirole[92], but without considering coalitional agreements.

The relationship between extensive form rationalizability and iterated deletion of weakly dominated strategies in the normal form representation is investigated in Marx and Swinkels[97] and Shimoji and Watson[98]. Related is the idea of normal form information sets of Mailath, Samuelson and Swinkels[94]. It is possible to extend the analysis of coalitional rationalizability to normal form information sets and obtain similar results to those in the above papers. In particular, it is straightforward to provide a notion of weak supported restriction to get a parallel concept with weakly dominated strategies (with strategies that are never best responses to conjectures that have full support over a certain set of strategies). For ease of exposition we give the definition for games of complete information.

For every $j \in N$ and $f_{-j} \in \Omega_{-j}$, let $\hat{u}_j(f_{-j}) \equiv u_j(s_j, f_{-j})$ for some $s_j \in BR_j(f_{-j})$.

Let A and B be such that $A \subset S$, $A = \times_{i \in I} A_i$ and $A \neq \emptyset$, $B \subset A$, $B = \times_{i \in I} B_i$ and $B \neq \emptyset$. Let $J \subset N$.

Definition: B is a *weakly supported restriction* by J given A if

- 1) $B_i = A_i, \forall i \notin J$, and
- 2) $\forall j \in J, f_{-j} \in \Omega_{-j}(A)$ such that there is $s_j \in A_j/B_j$ for which $s_j \in BR_j(f_{-j})$ it is the case that $\hat{u}_j(f_{-j}) \leq \hat{u}_j(g_{-j}) \forall g_{-j}$ such that $g_{-j} \in \Omega_{-j}(B)$ and $g_{-i}^{-J} = f_{-i}^{-J}$, and $\exists g_{-j} \in \Omega_{-j}(B)$ such that $g_{-i}^{-J} = f_{-i}^{-J}$ and $\hat{u}_j(f_{-j}) < \hat{u}_j(g_{-j})$.

A supported restriction is weakly supported by some coalition if, fixing the conjecture concerning strategies played by players outside the coalition, making the restriction is weakly preferred by every player in the coalition to playing a strategy that violates the restriction.

The well-known problems associated with iterated deletion of weakly dominated strategies apply equally to the iterative removal of strategies that are not part of some weakly supported restriction. In particular the set of strategies obtained by the procedure depends on the order in which restrictions are made.

An additional issue which arises is that the intersection of weakly supported restrictions can yield an empty set. Consider the game of Figure 15.

	C1		C2	
	B1	B2	B1	B2
A1	2,1,2	2,2,2	1,1,1	1,1,1
A2	2,2,1	1,1,1	2,2,1	1,1,1

Figure 15

In the above game $\{A1\} \times \{B1, B2\} \times \{C1\}$ is a weakly supported restriction by $\{1, 3\}$ given S . Similarly, $\{A2\} \times \{B1\} \times \{C1, C2\}$ is a weakly supported restriction by $\{1, 2\}$ given S . The intersection of these restrictions is the empty set because there is no strategy of player 1 that is inside both restrictions.

This example is in contrast with the result that the set of extensive form coalitionally rationalizable strategies is always nonempty. One could try to identify conditions on normal form games under which iterated removal of strategies that are not part of some weakly supported restriction is insensitive to the order of restrictions and leads to a nonempty set. Furthermore, one could look for conditions on extensive form games that guarantee the above conditions in the normal form representation of the game, and conditions which guarantee that the set of outcomes that are consistent with extensive form coalitional rationalizability is equivalent to the set of outcomes consistent with iterated removal of strategies that are not part of a weakly supported restriction in the normal form representation of the game. This project is left for future research.

9 Conclusion

This paper shows that coalitional and sequential rationality considerations can be incorporated simultaneously in a non-equilibrium framework in multi-stage games in a manner that existence is guaranteed, even in incomplete information environments. We consider this to be a step in understanding the structure of coalitional agreements in general extensive form games. An important future direction is investigating the case when coalitionally rational players to be able to communicate to each other during the play of an extensive form game. If this analysis involves making assumptions on how communication affects the beliefs of players' involved, one should also consider whether players have the option not to participate in such communication.

10 Appendix

Lemma 1: Let $(\underline{B}^1, \dots, \underline{B}^k, \underline{B}') \in \mathcal{M}$ and $(\underline{B}^1, \dots, \underline{B}^k, \underline{B}'') \in \mathcal{M}$. If $\underline{B}' \cap \underline{B}'' \neq \emptyset$ then $(\underline{B}^1, \dots, \underline{B}^k, \underline{B}' \cap \underline{B}'') \in \mathcal{M}$.

Proof: Take an arbitrary $i \in N$, $\tau_i \in \Upsilon_i$ and $\theta_{-i} \in \Theta_{-i}^c(\underline{B}^1, \dots, \underline{B}^k, \underline{B}' \cap \underline{B}'')$. Let $s_i \in BR_{\tau_i}(\theta_{-i})$. Since $\theta_{-i} \in \Theta_{-i}^c(\underline{B}^1, \dots, \underline{B}^k, \underline{B}')$ and $\theta_{-i} \in \Theta_{-i}^c(\underline{B}^1, \dots, \underline{B}^k, \underline{B}'')$, by the starting hypotheses $s_i \in \underline{B}'_{\tau_i}$ and $s_i \in \underline{B}''_{\tau_i}$ and hence $s_i \in \underline{B}'_{\tau_i} \cap \underline{B}''_{\tau_i}$. Since i and τ_i were arbitrary, this establishes the lemma. QED

Proof of Claim 1: For every $i \in N$, $\tau_i \in \Upsilon_i$ and $\omega_{-i} \in \Omega_{-i}$ define $\hat{u}_{\tau_i}(\omega_{-i}) \equiv \max_{s_i \in S_i} u_{\tau_i}(s_i, \omega_{-i})$. Fix some $i \in N$ and $\tau_i \in \Upsilon_i$. Construct $\theta_{-i}^{\tau_i} \in \Theta_{-i}$ the following way. Let $\theta_{-i}^{\tau_i}(\emptyset)$ be such that $\theta_{-i}^{\tau_i}(\emptyset) \in \Omega_{-i}(\underline{B}^k)$ and $\hat{u}_{\tau_i}(\theta_{-i}^{\tau_i}(\emptyset)) \geq \hat{u}_{\tau_i}(\omega_{-i}) \forall \omega_{-i} \in \Omega_{-i}(\underline{B}^k)$. Standard arguments establish that there exists such a belief $\theta_{-i}^{\tau_i}(\emptyset)$. Let $H_0^{\tau_i} \equiv \{h \in H/H^Z \mid h \text{ is reached by } \theta_{-i}^{\tau_i}(h_0)\}$. Let $\theta_{-i}^{\tau_i}(h) = \theta_{-i}^{\tau_i}(h_0) \forall h \in H_0^{\tau_i}$. Next, we give an iterative method to define beliefs at remaining histories. Assume $H_k^{\tau_i}$ is defined for some $k \geq 0$. Then let $I_{k+1}^{\tau_i} \equiv \{h \in H/H^Z \mid h \notin H_k^{\tau_i}, \text{ but } \text{imp}(h) \in H_k^{\tau_i}\}$. For any $h \in I_{k+1}^{\tau_i}$ define $\theta_{-i}^{\tau_i}(h)$ the following way. Let m (where $0 \leq m \leq k$) be the highest index such that \underline{B}^m reaches h . Then let $\theta_{-i}^{\tau_i}(h)$ be such that $\theta_{-i}^{\tau_i}(h)$ reaches h , $\theta_{-i}^{\tau_i}(h) \in \Omega_{-i}(\underline{B}^m)$ and $\hat{u}_{\tau_i}^h(\theta_{-i}^{\tau_i}(h)) \geq \hat{u}_{\tau_i}^h(\omega_{-i}) \forall \omega_{-i} \in \Omega_{-i}$ such that ω_{-i} reaches h and $\omega_{-i} \in \Omega_{-i}(\underline{B}^m)$. Again it is straightforward to establish the existence of such a belief $\theta_{-i}^{\tau_i}(h)$. Let $H_{k+1}^{\tau_i}(h) \equiv \{h' \in h \cup \text{succ}(h) \mid h' \text{ is reached by } \theta_{-i}^{\tau_i}(h)\}$. For all $h' \in H_{k+1}^{\tau_i}(h)$ let $\theta_{-i}^{\tau_i}(h') = \theta_{-i}^{\tau_i}(h)$. Finally, let $H_{k+1}^{\tau_i} \equiv \bigcup_{h \in I_{k+1}^{\tau_i}} H_{k+1}^{\tau_i}(h)$.

By construction $\theta_{-i}^{\tau_i} \in \Theta_{-i}^c(\underline{B}^1, \dots, \underline{B}^k)$. Let now $s_i \in BR_{\tau_i}(\theta_{-i}^{\tau_i})$. Since $(\underline{B}^1, \dots, \underline{B}^k) \in \mathcal{M}$, $s_i \in \underline{B}^k_{\tau_i}$. Suppose now that $s_i \notin C_{\tau_i}$. That means there is $h \in H/H^Z$, $\underline{B} \subset \underline{B}^k$ and $\Phi \subset \bigcup_{i \in N} \Upsilon_i$ such that \underline{B} is an h -based supported restriction by Φ given \underline{B}^k and $s_i \notin \underline{B}_{\tau_i}$. By the definition of the supported restriction then s_i reaches h . For every $j \in N/\{i\}$ and $v_j \in \Upsilon_j$ let $T_{v_j} : S_j \rightarrow S_j$ be such that for every $s_j \in \underline{B}_{v_j}^k / \underline{C}_{v_j}$ that reaches h it holds that $T_{v_j}(s_j) \in \underline{C}_{v_j}$ and $T_{v_j}(s_j)$ is an h -replacement of s_j , and for every other $s_j \in S_j$ it holds that $T_{v_j}(s_j) = s_j$. Let now $\psi_{-i}^{\tau_i} \in \Theta_{-i}$ be such that $\psi_{-i}^{\tau_i} \in \Theta_{-i}^c(\underline{B}^1, \dots, \underline{B}^k, \underline{C})$ and for every $h' \notin \text{succ}(h)$ it holds that $\psi_{-i}^{\tau_i}[h'](\underline{s}) = 0$ if there is no $\underline{t} \in \underline{S}$ for which $\underline{t}_{v_j} = T_{v_j}(\underline{s}_{v_j}) \forall j \in N/\{i\}$ and $v_j \in \Upsilon_j$ and let $\psi_{-i}^{\tau_i}[h'](\underline{s}) = \sum_{\underline{t} \in \underline{S}: \underline{t}_{v_j} = T_{v_j}(\underline{s}_{v_j}) \forall j \in N/\{i\}, v_j \in \Upsilon_j} \theta_{-i}^{\tau_i}[h'](\underline{t})$ otherwise. The existence of such a belief

system $\psi_{-i}^{\tau_i}$ is straightforward from $\theta_{-i}^{\tau_i} \in \Theta_{-i}^c(\underline{B}^1, \dots, \underline{B}^k)$ and the definition of T . Note that by construction $\theta_{-i}^{\tau_i}[\underline{h}] = \psi_{-i}^{\tau_i}[\underline{h}]$. Since $s_i \in BR_{\tau_i}(\theta_{-i}^{\tau_i})$ and $(\underline{B}^1, \dots, \underline{B}^k) \notin \mathcal{M}$, by the construction of $\psi_{-i}^{\tau_i}$ it follows that $\exists b_i \in \underline{B}_{\tau_i}^k$ such that $b_i(h') = s_i(h') \forall h' \notin h \cup \text{succ}(h)$. Since $\psi_{-i}^{\tau_i} \in \Theta_{-i}^c(\underline{B}^1, \dots, \underline{B}^k, \underline{B})$,

$b_i \in \underline{B}_{\tau_i}$ by the definition of a supported restriction. Then since \underline{B} is a supported restriction by Φ given \underline{B}^k , $u_{\tau_i}^h(\psi^{\tau_i}(h), b_i) > u_{\tau_i}^h(\theta_{-i}^{\tau_i}(h), s_i)$. But by construction $u_{\tau_i}^h(\theta_{-i}^{\tau_i}(h), s_i) \geq u_{\tau_i}^h(\psi^{\tau_i}(h), b_i)$, a contradiction! This implies that $s_i \in C_{\tau_i}$. Since τ_i was arbitrary, this implies $\underline{C} \neq \emptyset$.

Now suppose $\underline{B} \in F(\underline{B}^k)$ and $(\underline{B}^1, \dots, \underline{B}^k, \underline{B}) \notin \mathcal{M}$. Since $(\underline{B}^1, \dots, \underline{B}^k) \in \mathcal{M}$ this implies $\exists i \in N$, $\tau_i \in \Upsilon_i$, $s_i \in B_{\tau_i}^k / B_{\tau_i}$ and $\theta_{-i} \in \Theta_{-i}^c(\underline{B}^1, \dots, \underline{B}^k, \underline{B})$ such that $s_i \in BR_{\tau_i}(\theta_{-i})$. But that contradicts that \underline{B} is a supported restriction from \underline{B}^k by some coalition of types. Therefore $(\underline{B}^1, \dots, \underline{B}^k, \underline{B}) \in \mathcal{M} \forall \underline{B} \in F(\underline{B}^k)$. Then Lemma 1 implies that $(\underline{B}^1, \dots, \underline{B}^k, \bigcap_{\underline{B} \in F(\underline{B}^k)} \underline{B}) \in \mathcal{M}$. QED

Proof of Claim 2: The sequence $(\underline{A}^k)_{k=0,1,\dots}$ is decreasing, so finiteness of \underline{S} implies that $\exists K \geq 0$ such that $\underline{A}^* = \underline{A}^k \forall k \geq K$. Since the trivial sequence $(\underline{A}^0) = (\underline{S}) \in \mathcal{M}$, Claim 1 implies that $\underline{A}^k \neq \emptyset$ and $(\underline{A}^0, \dots, \underline{A}^k) \in \mathcal{M} \forall k \geq 0$. In particular then $\underline{A}^K = \underline{A}^* \neq \emptyset$ and $(\underline{A}^0, \dots, \underline{A}^K) \in \mathcal{M}$. QED

Proof of Claim 3: If \underline{B} is a restriction by $\{\tau_i\}$ given \underline{A} and for every $s_i \in A_{\tau_i} / B_{\tau_i}$ it holds that there is no $\theta_{-i} \in \Theta_{-i}^c(\underline{A})$ such that $s_i \in BR_{\tau_i}(\theta_{-i})$ then the definition of a supported restriction trivially implies that \underline{B} is a null-history based supported restriction. Suppose now that \underline{B} is a h -based supported restriction by $\{\tau_i\}$ for some $h \in H/H^Z$. Then the definition of supported restriction implies that no $s_i \in A_{\tau_i} / B_{\tau_i}$ can be a best response among h -replacement strategies to any $\theta_{-i} \in \Theta_{-i}^c(\underline{A})$. Therefore $s_i \notin BR_{\tau_i}(\theta_{-i})$ for every $s_i \in B_{\tau_i}^k / B_{\tau_i}$ and $\theta_{-i} \in \Theta_{-i}^c(\underline{A})$. QED

Lemma 2: For every $k \geq 1$, $i \in N$, $\tau_i \in \Upsilon_i$ and $s_i \in A_{\tau_i}^k$ it holds that there is $\theta_{-i} \in \Theta_{-i}^c(\underline{A}^0, \dots, \underline{A}^{k-1})$ such that $s_i \in BR_{\tau_i}(\theta_{-i})$.

Proof: Claim 3 implies that the claim holds for $k = 1$. Suppose now it holds for some $k \geq 1$. Then for every $i \in N$, $\tau_i \in \Upsilon_i$ and $s_i \in A_{\tau_i}^k$ it holds that there is $\theta_{-i} \in \Theta_{-i}^c(\underline{A}^0, \dots, \underline{A}^{k-1})$ such that $s_i \in BR_{\tau_i}(\theta_{-i})$. Furthermore Claim 3 implies that there is $\theta'_{-i} \in \Theta_{-i}^c(\underline{A}^k)$ such that $s_i \in BR_{\tau_i}(\theta'_{-i})$. Now construct $\theta''_{-i} \in \Theta_{-i}$ such that $\theta''_{-i}(h) = \theta'_{-i}(h) \forall h$ such that h is reached by \underline{A}^k_{-i} . At any other history h , let $\theta''_{-i}(h) = \theta_{-i}(h)$. Belief process θ''_{-i} is consistent because θ_{-i} and θ'_{-i} are consistent and because $\theta'_{-i} \in \Theta_{-i}^c(\underline{A}^k)$ implies that there can be no $h \in H/H^Z$ such that there is h' reached by \underline{A}^k_{-i} and $\theta'_{-i}(h')$ reaches h . By construction θ''_{-i} is concentrated on $(\underline{A}^0, \dots, \underline{A}^k)$. Finally, $s_i \in BR_{\tau_i}(\theta''_{-i})$ since at every $h \in H/H^Z$ either $\theta''_{-i}(h) = \theta_{-i}(h)$ or $\theta''_{-i}(h) = \theta'_{-i}(h)$, so $s_i \in BR_{\tau_i}(\theta_{-i})$ and $s_i \in BR_{\tau_i}(\theta'_{-i})$ imply that s_i is a best response among h -replacements to $\theta''_{-i}(h)$ at every $h \in H/H^Z$. QED

Proof of Claim 4: By Claim 2 $(\underline{A}^0, \dots, \underline{A}^K) \in \mathcal{M}$. Furthermore by Claim 2 there is no nontrivial supported restriction by any coalition given \underline{A}^K . Claim 3 then implies that for every $i \in N$, $\tau_i \in \Upsilon_i$ and $s_i \in A_{\tau_i}^K$ there is $\theta_{-i}^K \in \Theta_{-i}^c(\underline{A}^K)$ such that $s_i \in BR_{\tau_i}(\theta_{-i}^K)$. Furthermore, note that $s_i \in A_{\tau_i}^k \forall k = 1, \dots, K$ and then Claim 3 and the definition of the sequence $(\underline{A}^0, \dots, \underline{A}^K)$ implies that for

every $k = 0, 1, \dots, K-1 \exists \theta_{-i}^k \in \Theta_{-i}^c(\underline{A}^k)$ such that $s_i \in BR_{\tau_i}(\theta_{-i}^k)$. Let now $\widehat{\theta}_{-i} \in \Theta_{-i}$ be such that $\widehat{\theta}_{-i}[h] = \theta_{-i}^K[h]$ if h is reached by \underline{A}_{-i}^K , and for every $k = 0, \dots, K-1 \widehat{\theta}_{-i}[h] = \theta_{-i}^k[h]$ if h is reached by \underline{A}_{-i}^k but not reached by \underline{A}_{-i}^{k+1} . By construction $\widehat{\theta}_{-i} \in \Theta_{-i}^c(\underline{A}^0, \dots, \underline{A}^K)$ and $s_i \in BR_{\tau_i}(\widehat{\theta}_{-i})$, which establishes the claim. QED

Proof of Claim 5: Essentially the same proofs as for the ones for Claims 1 through 4 establish that there is $L \geq 0$ such that $\underline{B}^k = \underline{B}^L, \forall k \geq L$, that $\underline{B}^L \neq \emptyset$, that $(\underline{B}^0, \dots, \underline{B}^L) \in \mathcal{M} \forall l = 0, 1, 2, \dots$ and that $(\underline{B}^0, \dots, \underline{B}^L)$ is coherent.

Define $H^{\tau_i} = \{h \in H/H^Z \mid \exists t \in S \text{ reaching } h \text{ and } v_{-i} \in \Upsilon_{-i} \text{ for which } t_i \in \underline{A}_{\tau_i}^* \text{ and } \forall j \in N/\{i\} t_j \in \underline{A}_{v_j}^* \text{ and } \varphi_i(\tau_i)(v_{-i}) > 0\}$.

Suppose that $k \geq 0$ is such that $\forall i \in N, \tau_i \in \Upsilon_i$ and $s_i \in \underline{A}_{\tau_i}^* \exists s'_i \in \underline{B}_{\tau_i}^k$ such that $s'_i(h) = s_i(h) \forall h \in H^{\tau_i}$. Note that this implies that $O_\tau(\underline{B}^k) \supset O_\tau(\underline{A}^*) \forall \tau \in \Upsilon$ for which $\varphi_i(\tau_i)(\tau_{-i}) > 0 \forall i \in N$.

Suppose now that $\exists h \in H/H^Z$ and $\underline{C} \subset \underline{B}^k$ such that \underline{C} is an h -based supported restriction by Φ given \underline{B}^k and that $\exists i \in N, \tau_i \in \Upsilon_i$ and $s_i \in \underline{A}_{\tau_i}^*$ such that there is no $s'_i \in \underline{B}_{\tau_i}^k$ for which $s'_i(h) = s_i(h) \forall h \in H^{\tau_i}$.

By the definition of supported restriction it has to be that $h \in H^{\tau_i}$.

For every $i \in N, \tau_i \in \Upsilon_i$ let $S_{\tau_i}^k = \{t_i \in \underline{A}_{\tau_i}^* \mid \exists t'_i \in \underline{B}_{\tau_i}^k \text{ for which } t'_i(h) = t_i(h) \forall h \in H^{\tau_i}\}$ and let $\Phi' = \bigcup_{i \in N} \{\tau_i \in \Upsilon_i \cap \Phi \mid S_{\tau_i}^k \neq \underline{A}_{\tau_i}^*\}$.

Suppose first that $\forall i \in N, \tau_i \in \Upsilon_i$ it holds that $S_{\tau_i}^k \neq \emptyset$. Let $\underline{S}^k = \prod_{i \in N} S_{\tau_i}^k$ and consider the restriction \underline{S}^k by Φ' given \underline{A}^* . By construction this restriction is h -based. Let $i \in N, \tau_i \in \Upsilon_i \cap \Phi', \theta_{-i} \in \Theta_{-i}^c(\underline{A}^*)$ and $s_i \in \underline{A}_{\tau_i}^*/\underline{S}_{\tau_i}^k$ be such that s_i reaches h and $s_i \in BR_{\tau_i}(\theta_{-i})$. Let t_i, ψ_{-i} such that $\psi_{-i} \in \Theta_{-i}^c(\underline{S}^k), \theta_{-\Phi'}(h) = \psi_{-\Phi'}(h), t_i \in \underline{A}_{\tau_i}$ and t_i is a best response to ψ_{-i} among h' -replacement strategies of s_i for every $h' \in h \cup \text{succ}(h)$ that is reached by t_i . For every $i \in N$ and $\tau_i \in \Upsilon_i$ let $T_{\tau_i} : S_i \rightarrow S_i$ such that for every $a_i \in \underline{A}_{\tau_i}^*$ it holds that $T_{\tau_i}(a_i) \in \underline{B}_{\tau_i}^k$ and $T_{\tau_i}(a_i)(h') = a_i(h') \forall h' \in H^{\tau_i}$ and for every $b_i \notin \underline{A}_{\tau_i}^*$ it holds that $T_{\tau_i}(b_i) = b_i$. By the starting assumption above there is a function T_{τ_i} like this. Let now θ'_{-i} such that $\theta'_{-i} \in \Theta_{-i}^c(\underline{B}^1, \dots, \underline{B}^k)$ and $\forall h' \in H^{\tau_i}$ it holds that $\theta'_{-i}[h'](b_i) = 0$ if there is no $a_i \in \underline{A}_{\tau_i}^*$ for which $T_{\tau_i}(a_i) = b_i$ and let $\theta'_{-i}[h'](b_i) = \sum_{t_i: T_{\tau_i}(t_i)=b_i} \theta_{-i}[h'](a_i)$ otherwise. The construction of such

belief process θ'_{-i} is tedious, but straightforward. Similarly let ψ'_{-i} such that $\psi'_{-i} \in \Theta_{-i}^c(\underline{B}^1, \dots, \underline{B}^k, \underline{C})$ and $\forall h' \in H^{\tau_i}$ it holds that $\theta'_{-i}[h'](b_i) = 0$ if there is no $a_i \in \underline{A}_{\tau_i}^*$ for which $T_{\tau_i}(a_i) = b_i$ and let $\psi'_{-i}[h'](b_i) = \sum_{t_i: T_{\tau_i}(t_i)=b_i} \psi_{-i}[h'](a_i)$

otherwise. Since $\psi_{-i} \in \Theta_{-i}^c(\underline{S}^k), \exists \psi'_{-i} \in \Theta_{-i}^c(\underline{B}^1, \dots, \underline{B}^k, \underline{C}^r)$ like this. Note that by construction $\theta_{-\Upsilon^r \cap \Phi}(h) = \psi_{-\Upsilon^r \cap \Phi}(h)$. Since $(\underline{B}^1, \dots, \underline{B}^L) \in \mathcal{M}, \exists s'_i \in BR_{\tau_i}(\theta'_{-i})$ such that $s'_i(h') = s_i(h') \forall h' \in H^{\tau_i}$. By construction $s'_i \notin \underline{C}_{\tau_i}$. Similarly, $\exists t'_i \in BR_{\tau_i}(\theta'_{-i})$ such that $t'_i(h') = t_i(h') \forall h' \in H^{\tau_i}$. But then the fact that \underline{C} is an h -based supported restriction by Φ given \underline{B}^k implies that

$u_{\tau_i}^h(t'_i, \theta'_{-i}) < u_{\tau_i}^h(s'_i, \psi'_{-i})$. Since $u_{\tau_i}^h(t'_i, \theta'_{-i}) = u_{\tau_i}^h(t_i, \theta_{-i})$ and $u_{\tau_i}^h(s'_i, \psi'_{-i}) = u_{\tau_i}^h(s_i, \psi_{-i})$, this implies $u_{\tau_i}^h(t_i, \theta_{-i}) < u_{\tau_i}^h(s_i, \psi_{-i})$. But that means \underline{S}^k is a nontrivial supported restriction by Φ' given \underline{A}^* , contradicting that there is no nontrivial supported restriction given \underline{A}^* .

Suppose now that for some $\forall i \in N$, $\tau_i \in \Upsilon_i$ it holds that $S_{\tau_i}^k = \emptyset$. That implies $\underline{A}^* \cap \underline{C} = \emptyset$. Let now $l \geq 0$ be such that $\forall i \in N$, $\tau_i \in \Upsilon_i$ there is $s_i \in \underline{A}_{\tau_i}^l \cap \underline{C}_{\tau_i}$ that reaches h , while there is some $i \in N$, $\tau_i \in \Upsilon_i$ for which

there is no $s_i \in \underline{A}_{\tau_i}^{l+1} \cap \underline{C}_{\tau_i}$ that reaches h .

For every $v_j \in \Upsilon_j$ let $\widehat{\underline{C}}_{v_i} = \{s_i \in \underline{C}_{v_i} \mid s_i \text{ reaches } h\}$. Assume that for $l \geq 0$ it holds that $\widehat{\underline{C}}_{v_i} \cap \underline{A}_{v_i}^l \neq \emptyset$. Let $i \in N$ and $\tau_i \in \Upsilon_i$ such that $h \in H^{\tau_i}$. Let $\theta_{-i} \in \Theta_{-i}^c(\underline{A}^1, \dots, \underline{A}^K)$ and $a_i \in \underline{A}_{\tau_i}^*$ be such that $a_i \in BR_{\tau_i}(\theta_{-i})$ and a_i reaches h . The existence of such θ_{-i} and a_i is guaranteed by $(\underline{A}^1, \dots, \underline{A}^K)$ being coherent and because $h \in H^{\tau_i}$. For every $j \in N/\{i\}$ and $v_j \in \Upsilon_j$ define the collection of functions \mathcal{T}_{v_j} such that $T_{v_j} \in \mathcal{T}_{v_j}$ iff $T_{v_j} : S_j \rightarrow S_j$, and for every $s_j \in \underline{A}_{v_j}^*$ that reaches h it holds that $T_{v_j}(s_j) \in \widehat{\underline{C}}_{v_i} \cap \underline{A}_{v_i}^l$ and $T_{v_j}(s_j)$ is an h -replacement of s_j , and for every other $s_j \in S_j$ it holds that $T_{v_j}(s_j) = s_j$. The starting assumption on \underline{A}^l guarantees that \mathcal{T}_{v_j} is nonempty. For every $T = \prod_{j \in N/\{i\}} \times_{v_j \in \Upsilon_j} T_{v_j} \in \prod_{j \in N/\{i\}} \times_{v_j \in \Upsilon_j} \mathcal{T}_{v_j}$ let $\widehat{\theta}_{-i}^T$ be such that $\widehat{\theta}_{-i}^T[h'](\underline{s}) = \sum_{\underline{t} \in \underline{S} : \underline{t}_{v_j} = T_{v_j}(\underline{t}_{v_j}) \forall j \in N/\{i\} \text{ and } v_j \in \Upsilon_j} \theta_{-i}[h'](\underline{t}) \forall h' \notin \text{succ}(h)$. For $h' \in \text{succ}(h)$

define $\widehat{\theta}_{-i}^T[h']$ iteratively such that $\widehat{\theta}_{-i}^T[h'] = \widehat{\theta}_{-i}^T[\text{imp}(h')]$ if $\widehat{\theta}_{-i}^T[\text{imp}(h')]$ reaches h' , otherwise such that $\widehat{\theta}_{-i}^T[h'] \in \Omega_{-i}(\underline{C} \cap \underline{A}^l)$ if $\underline{C}_{-i} \cap \underline{A}_{-i}^l$ reaches h' , for every $k = 1, \dots, l$ $\widehat{\theta}_{-i}^T[h'] \in \Omega_{-i}(\underline{A}^k)$ if \underline{A}_{-i}^k reaches h' , and there is $a_i \in \underline{A}_{\tau_i}^l$ such that $u_{\tau_i}^h(a_i, \widehat{\theta}_{-i}^T) \geq u_{\tau_i}(b_i, \omega_{-i}) \forall b_i \in \underline{A}_{\tau_i}^l$ and $\omega_{-i} \in \Omega_{-i}$ such that $\omega_{-i} \in \Omega_{-i}(\underline{C} \cap \underline{A}^l)$ if $\underline{C}_{-i} \cap \underline{A}_{-i}^l$ reaches h' , and for every $k = 1, \dots, l$ $\omega_{-i} \in \Omega_{-i}(\underline{A}^k)$ if \underline{A}_{-i}^k reaches h' . It is straightforward to establish the existence of such belief process $\widehat{\theta}_{-i}^T$. By construction $\widehat{\theta}_{-i}^T \in \Theta_{-i}^c(\underline{A}^1, \dots, \underline{A}^l, \widehat{\underline{C}})$. Since $\Theta_{-i}^c(\underline{A}^1, \dots, \underline{A}^l, \widehat{\underline{C}}) \in \mathcal{M}$, for any $c_i \in BR_{\tau_i}(\widehat{\theta}_{-i}^T)$ it holds that $c_i \in \underline{C}_{\tau_i}$. Since $a_i \in BR_{\tau_i}(\theta_{-i})$ and the way $\widehat{\theta}_{-i}^T$ is created from θ_{-i} , there is $c_i^* \in BR_{\tau_i}(\widehat{\theta}_{-i}^T)$ such that $c_i^*(h') = a_i(h') \forall h' \notin h \cup \text{succ}(h)$. Since a_i reaches h , this implies that c_i^* reaches h . Now for every $T \in \prod_{j \in N/\{i\}} \times_{v_j \in \Upsilon_j} \mathcal{T}_{v_j}$ let $c_i^*(T) \in BR_{\tau_i}(\widehat{\theta}_{-i}^T)$ be such that $c_i^*(T)$ reaches h .

Let $T^* = \prod_{j \in N/\{i\}} \times_{v_j \in \Upsilon_j} \mathcal{T}_{v_j}$ be such that $u_{\tau_i}^h(c_i^*(T^*), \widehat{\theta}_{-i}^{T^*}) \geq u_{\tau_i}^h(c_i^*(T), \widehat{\theta}_{-i}^T) \forall T \in \prod_{j \in N/\{i\}} \times_{v_j \in \Upsilon_j} \mathcal{T}_{v_j}$. The existence of such T^* is guaranteed by the finiteness of

the type space. Since $\underline{C} \cap \underline{A}^* = \emptyset$ by the starting assumption, $u_{\tau_i}^h(c_i^*(T^*), \widehat{\theta}_{-i}^{T^*}) > u_{\tau_i}^h(a_i, \theta_{-i})$. Since $c_i^*(T^*)(h') = a_i(h') \forall h' \notin h \cup \text{succ}(h)$ and $a_i \in \underline{A}_{\tau_i}^*$, this implies for any $h' \notin h \cup \text{succ}(h)$ there is no h' -based supported restriction

\underline{D} given \underline{A}^l such that $c_i^*(T^*) \notin \underline{D}_{\tau_i}$. Similar arguments as in the proof of Claim 1 establish that the construction of $c_i^*(T^*)$ and $\widehat{\theta}_{-i}^{T^*}$ imply that for any $h' \notin h \cup \text{succ}(h)$ there is no h' -based supported restriction \underline{D} given \underline{A}^l such that $c_i^*(T^*) \notin \underline{D}_{\tau_i}$. This establishes that $c_i^*(T^*) \in \underline{A}_{\tau_i}^{l+1}$. This concludes that for every $i \in N$ and $\tau_i \in \Upsilon_i$ for which $h \in H^{\tau_i}$ it holds that $\widehat{C}_{\tau_i} \cap \underline{A}_{\tau_i}^{l+1} \neq \emptyset$. But then since $\widehat{C}_{\tau_i} \cap \underline{A}_{\tau_i}^0 = \widehat{C}_{\tau_i} \neq \emptyset \vee \widehat{C}_{\tau_i}^r$, it follows that $\widehat{C}_{\tau_i} \cap \underline{A}_{\tau_i}^* \neq \emptyset$. But this contradicts the starting hypothesis that for some $\forall i \in N, \tau_i \in \Upsilon_i$ it holds that $S_{\tau_i}^k = \emptyset$.

This concludes that $\forall i \in N, \tau_i \in \Upsilon_i$ and $s_i \in \underline{A}_{\tau_i}^* \exists s'_i \in \underline{D}_{\tau_i}^{k+1}$ such that $s'_i(h) = s_i(h) \forall h \in H^{\tau_i}$. Since $\forall i \in N, \tau_i \in \Upsilon_i$ and $s_i \in \underline{A}_{\tau_i}^* \exists s'_i \in \underline{B}_{\tau_i}^0$ such that $s'_i(h) = s_i(h) \forall h \in H^{\tau_i}$, by induction $\forall i \in N, \tau_i \in \Upsilon_i$ and $s_i \in \underline{A}_{\tau_i}^* \exists s'_i \in \underline{B}_{\tau_i}^L$ such that $s'_i(h) = s_i(h) \forall h \in H^{\tau_i}$. This implies that $O_\tau(\underline{B}^L) \supset O_\tau(\underline{A}^*) \forall \tau \in \Upsilon$ for which $\varphi_i(\tau_i)(\tau_{-i}) > 0 \forall i \in N$.

A symmetric argument to the above establishes that $O_\tau(\underline{B}^L) \subset O_\tau(\underline{A}^*) \forall \tau \in \Upsilon$ for which $\varphi_i(\tau_i)(\tau_{-i}) > 0 \forall i \in N$. QED

Proof of Claim 6: Let $K' \geq 0$ be such that $\underline{R}^k \neq \underline{R}^{k-1} \forall k \leq K$ and $\underline{R}^k = \underline{R}^* \forall k \geq K'$. It is easy to establish the existence of such a K from the finiteness of the game. Let $\underline{B}^k = \underline{R}^k$ for $k = 0, \dots, K'$. For $k > K'$ define \underline{B}^k iteratively such that $\underline{B}^k = \bigcap_{\underline{B} \in F(\underline{B}^{k-1})} \underline{B}$. Note that $\underline{B}^0 = \underline{S}$. By Claim 3 $(\underline{B}^k)_{k=0,1,2,\dots}$ is such that for any $k \geq 0$ $\underline{B}^{k+1} = \underline{B}^k$ if there is no nontrivial supported restriction given \underline{B}^k , otherwise there is some Θ^k nonempty collection of nontrivial restrictions given \underline{B}^k for which $\underline{B}^{k+1} = \bigcap_{\underline{B} \in \Theta^k} \underline{B}$. Let $\tau \in \Upsilon$ be such that $\varphi_i(\tau_i)(\tau_{-i}) > 0 \forall i \in N$. Then by Claim 5 there is L such that $\underline{B}^k = \underline{B}^L \forall k \geq L$ and $O_\tau(\underline{B}^L) = O_\tau(\underline{A}^*)$. Since $\underline{B}^0, \underline{B}^1, \dots$ is a decreasing sequence, this implies $O_\tau(\underline{A}^*) = O(\underline{B}^L) \subset O(\underline{B}^{K'}) = O(\underline{R}^*)$. QED

Proof of Claim 7: If player 1 expects player 2 to play $B1$ in the first stage with more than probability $1 - \frac{1000}{10201} \approx 0.902$, then his best responses have to specify $A1$ in the first stage. To see this, note that player 1's expected first stage payoff is $100x + 9(1-x)$ when playing $A1$ and $10(1-x)$ when playing $A2$, where x is the probability that he expects player 1 to play $B1$ in the first stage. His minimal expected payoff in the second stage is $\frac{1000}{101}$ (when player 2 plays $B1$ with probability $\frac{1}{101}$, making player 1 indifferent between $A1$ and $A2$). His maximal expected payoff in the second stage is 100. So if $100x + 9(1-x) + \frac{1000}{101} > 10(1-x) + 100$, which is exactly when $x > 1 - \frac{1000}{10201}$, player 1 is always better off playing $A1$ in the first round, no matter what his conjectures are concerning continuation strategies of player 2. This implies that player 1 cannot be surprised after $(A2, B2)$ and $(A3, B2)$, because he can only play $A2$ or $A3$ if he expects player 1 to play $B2$ with positive probability. Next note that if player 1's strategy specifies playing $A2$ with positive probability

in stage 2 after $(A2, B2)$ in the first stage, then playing $A2$ in the first stage cannot be a best response if $x > (11 - \frac{1000}{101})/11 \approx 0.1$. This is because his payoff is at least $100x + 9(1-x) + \frac{1000}{101}$ if playing $A1$, while his payoff is at most $100x + 20(1-x)$ if he plays $A2$ in the first stage and then $A2$ in the second stage after $(A2, B2)$ in the first stage. Also note that if player 1's strategy specifies playing $A2$ with positive probability in stage 2 after $(A3, B2)$ in the first stage, then playing $A1$ in the first stage is always (for any conjecture) better for him than playing $A3$. This is because $100x + 9(1-x) + \frac{1000}{101}$ is strictly higher for every $x \in [0, 1]$ than $1 + 100x + 10(1-x)$, which is player 1's highest possible payoff if he plays $A3$ in the first period and then $A2$ in the second stage after $(A3, B2)$ in the first one. Therefore it is only rational for player 1 to play $A3$ if he plans to play $A1$ with probability 1 after $(A3, B2)$. But then a strategy in which player 2 plays $B2$ with positive probability can only be extensive form rationalizable if he plays $B1$ after $(A3, B2)$. So in any rationalizable outcome $(A1, B1)$ has to follow $(A3, B2)$. But then whenever player 1's strategy specifies playing $A2$ with positive probability in stage 2 after $(A2, B2)$ in the first stage, it is better for him to play $A3$ than $A2$ in the first stage if $x < \frac{81}{180 - \frac{1000}{101}} \approx 0.476$. This is because by the previous observation his payoff is at least $1 + \frac{1000}{101}x + 100(1-x)$ if he plays $A3$, while his payoff is at most $20(1-x) + 100x$ if he plays $A2$ in the first stage and then $A2$ in the second stage after $(A2, B2)$ in the first stage. Combining the above observations yields that it can only be extensive form rationalizable for player 1 to play $A2$ in the first stage if he plans to play $A1$ with probability 1 after the outcome $(A2, B2)$. But then it can only be extensive form rationalizable for player 2 to play $B2$ with positive probability in the first stage if he plans to play $B1$ in the second stage after $(A2, B2)$ in the first stage. This establishes that in any extensive form rationalizable (and therefore in every extensive form coalitionally rationalizable) outcome $(A1, B1)$ has to follow in the second stage the play of $(A2, B2)$ in the first stage. QED

Proof of Claim 8: For every $i \in N$, $\tau_i \in \Upsilon_i$ and $\underline{A} = \prod_{i \in N} (\prod_{\tau_i \in \Upsilon_i} \underline{A}_{\tau_i}) \subset \underline{S}$, $h \in H/H^Z$ reached by \underline{A}_{-i} , and $v_{-i} \in \Upsilon_{-i}$ let $\underline{S}_{-i}^{v_{-i}}(\underline{A}_{-i}, v_{-i}) = \{\underline{s}_{-i} \in \underline{S}_{-i} \mid \underline{s}_{v_1} \notin \underline{A}_{v_1} \text{ for some } j \neq i \text{ and } \underline{s}_{v_1} \text{ reaches } h \forall j \neq i\}$ and let $\underline{S}_{-i}^h(\underline{A}_{-i}, v_{-i}) = \{\underline{s}_{-i} \in \underline{S}_{-i} \mid \underline{s}_{v_1} \text{ reaches } h \forall j \neq i\}$. Then for every $\varepsilon \in (0, 1)$ let
$$\hat{\Theta}_{-\tau_i}(\underline{A}, \varepsilon) = \{\theta_{-i} \in \Theta_{-i}^c \mid \left(\frac{\sum_{v_{-i} \in \Upsilon_{-i}} \varphi(\tau_i)(v_{-i}) \sum_{\underline{s}_{-i} \in \underline{S}_{-i}^{v_{-i}}(\underline{A}_{-i}, v_{-i})} \theta_{-i}[h](\underline{s}_{-i})}{\sum_{v_{-i} \in \Upsilon_{-i}} \varphi(\tau_i)(v_{-i}) \sum_{\underline{s}_{-i} \in \underline{S}_{-i}^h(\underline{A}_{-i}, v_{-i})} \theta_{-i}[h](\underline{s}_{-i})} \right) < \varepsilon \forall h \in H/H^Z \text{ such that } \exists v_{-i} \in \Upsilon_{-i} \text{ and } \underline{s} \in \underline{S} \text{ such that } \underline{s}_{v_1} \text{ reaches } h \forall j \neq i\}.$$

From the definition of a supported restriction if for some $k \geq 0$ \underline{B} is an h -based nontrivial supported restriction by Φ given \underline{A}^k and $i \in N$, $\tau_i \in \Phi$ are such that $\underline{A}_{\tau_i}^k / \underline{B}_{\tau_i} \neq \emptyset$ and $b_i \in \underline{B}_{\tau_i}$, and $\theta_{-i} \in \Theta_{-i}^c(\underline{B})$ are such that $b_i \in BR_{\tau_i}(\theta_{-i})$ and b_i reaches h , then $u_{\tau_i}^h(b_i, \theta_{-i}) > \max_{a_i \in \underline{A}_{\tau_i}^k / \underline{B}_{\tau_i}} u_{\tau_i}^h(a_i, \theta_{-i})$.

Since $u_{\tau_i}^h$ is continuous in θ_{-i} (with respect to the weak topology) there is $\varepsilon(\underline{A}^k, \underline{B}, \Phi) > 0$ such that if $i \in N$, $\tau_i \in \Phi$ are such that $\underline{A}_{\tau_i}^k / \underline{B}_{\tau_i} \neq \emptyset$ and $b_i \in \underline{B}_{\tau_i}$ and $\theta_{-i} \in \widehat{\Theta}_{-i}(\underline{B}, \varepsilon)$ are such that $b_i \in BR_{\tau_i}(\theta_{-i})$ and b_i reaches h , then $u_{\tau_i}^h(b_i, \theta_{-i}) > \max_{a_i \in \underline{A}_{\tau_i}^k / \underline{B}_{\tau_i}} u_{\tau_i}^h(a_i, \theta_{-i})$. Since there are a finite number

of nontrivial supported restrictions given \underline{A}^k for any $0 \leq k \leq K-1$, there is $\widehat{\varepsilon} > 0$ such that for every $0 \leq k \leq K-1$, $i \in N$, $\tau_i \in \Phi$, $\underline{B} \subset \underline{A}^k$, $h \in H/H^Z$, $\theta_{-i} \in \widehat{\Theta}_{-i}(\underline{B}, \widehat{\varepsilon})$ and $b_i \in BR_{\tau_i}(\theta_{-i})$ for which \underline{B} is an h -based nontrivial supported restriction by Φ given \underline{A}^k , $\underline{A}_{\tau_i}^k / \underline{B}_{\tau_i} \neq \emptyset$ and b_i reaches h it holds that $u_{\tau_i}^h(b_i, \theta_{-i}) > \max_{a_i \in \underline{A}_{\tau_i}^k / \underline{B}_{\tau_i}} u_{\tau_i}^h(a_i, \theta_{-i})$.

Let now $\underline{\varepsilon}^1, \underline{\varepsilon}^2, \dots$ be such that the following three conditions hold:

- (i) for every $k \geq 1$ $\underline{\varepsilon}^k = (\underline{\varepsilon}_{\tau_i, s_i}^k)_{i \in N, \tau_i \in \Upsilon_i, s_i \in S_i}$ such that $\underline{\varepsilon}_{\tau_i, s_i}^k > 0 \forall i \in N, \tau_i \in \Upsilon_i, s_i \in S_i$ and $\sum_{i \in N, \tau_i \in \Upsilon_i, s_i \in S_i} \underline{\varepsilon}_{\tau_i, s_i}^k < 1$
- (ii) $\forall i \in N, \tau_i \in \Upsilon_i, s_i \in S_i$ it holds that $\underline{\varepsilon}_{\tau_i, s_i}^k \rightarrow 0$ as $k \rightarrow \infty$

(iii) $\forall i \in N, \tau_i \in \Upsilon_i, k \geq 1$ and $1 \leq m \leq K$ it holds that $\underline{\varepsilon}_{\tau_i, s_i}^k > \widehat{\varepsilon} \cdot \underline{\varepsilon}_{\tau_i, s_i'}^m$ if $s_i \in \underline{A}_{\tau_i}^m$ and $s_i' \notin \underline{A}_{\tau_i}^m$.

For every $k \geq 1$ construct the following perturbed game G^k from G . Let the set of players be Υ . For every $i \in N$ and $\tau_i \in \Upsilon_i$ let the set of strategies of τ_i in G^k be $\Sigma_{\tau_i}^k = \{\sigma_i \in \Delta(S_i) \mid \sigma_i(s_i) \geq \underline{\varepsilon}_{\tau_i, s_i}^k \forall s_i \in S_i \text{ and } \sigma_i(s_i) = \underline{\varepsilon}_{\tau_i, s_i}^k \forall s_i \notin \underline{A}_{\tau_i}^*\}$. It is the set of mixed strategies of τ_i that allocate exactly probability $\underline{\varepsilon}_{\tau_i, s_i}^k$ to every strategy $s_i \notin \underline{A}_{\tau_i}^*$ and at least probability $\underline{\varepsilon}_{\tau_i, s_i}^k$ to every strategy $s_i \in \underline{A}_{\tau_i}^*$. Let the payoff function of τ_i in G^k , denoted by $u_{\tau_i}^k$, be the relevant restriction of the mixed strategy extension of $u_{\tau_i}(\underline{s})$. Let $\underline{\sigma}$ denote a typical strategy profile in G^k .

Note that σ_{τ_i} is strictly mixed for every $i \in N, \tau_i \in \Upsilon_i, k \geq 1$ and $\sigma_{\tau_i} \in \Sigma_{\tau_i}^k$ and therefore for every $\underline{\sigma} \in \Sigma^k$ and for every $i \in N$, the belief process θ_{-i} which is defined by $\theta_{-i}[h] = \underline{\sigma}_{-i} \forall h \in H/H^Z$ is consistent. Denote this belief process for every $\underline{\sigma} \in \Sigma^k$ and $i \in N$ by $\theta_{-i}^{\underline{\sigma}}$. Note that every $\theta_{-i}^{\underline{\sigma}}$ defines an assessment (a probability distribution over decision nodes at every information set) for every $\tau_i \in \Upsilon_i$ through the beliefs over types $\varphi_i(\tau_i)$. Denote this assessment by $a_{\tau_i}^{\underline{\sigma}}$.

Since strategy sets are compact, convex and nonempty and payoff functions are linear, by Kakutani's fixed point theorem every G^k has a Nash equilibrium $\underline{\sigma}^{k*}$. Since $\underline{\sigma}^{k*}$ is strictly mixed, $\underline{\sigma}_{\tau_i}^{k*}$ is a (strong) sequential best response to $\underline{\sigma}_{-\tau_i}^{k*} \forall k \geq 1, i \in N$ and $\tau_i \in \Upsilon_i$. Therefore $\underline{\sigma}_{\tau_i}^{k*}$ is a best response to assessment $a_{\tau_i}^{\underline{\sigma}^{k*}}$. Since $\Delta(S_i)$ is compact, $(\underline{\sigma}^{k*})_{k=1,2,\dots}$ has a convergent subsequence. Then without loss of generality assume that $(\underline{\sigma}^{k*})_{k=1,2,\dots}$ is convergent, and let $\underline{\sigma}^* = \lim_{k \rightarrow \infty} \underline{\sigma}^{k*}$. Then $a_{\tau_i}^{\underline{\sigma}^{k*}}$ is convergent too for every $i \in N$ and $\tau_i \in \Upsilon_i$. Let $a_{\tau_i}^* = \lim_{k \rightarrow \infty} a_{\tau_i}^{\underline{\sigma}^{k*}}$. Note that $a_{\tau_i}^*$ is consistent with $\underline{\sigma}^*$. Also note that by construction $\text{supp} \underline{\sigma}_{\tau_i}^* \subset \underline{A}_{\tau_i}^* \forall i \in N$ and $\tau_i \in \Upsilon_i$.

By construction $\theta_{-i}^{\underline{\sigma}^*} \in \widehat{\Theta}_{-i}(\underline{A}^m, \widehat{\varepsilon})$ for every $i \in N, k \geq 1, \underline{\sigma} \in \Sigma^k$ and $1 \leq m \leq K$. Therefore if $s_i \in BR_{\tau_i}(\theta_{-i}^{\underline{\sigma}^*})$ in G , then $s_i \in \underline{A}_{\tau_i}^*$. This implies that

$\sigma_i \in BR_{\tau_i}(\theta_{-i}^{\underline{\sigma}})$ in G^k , then σ_i is not only a sequential best response to $\theta_{-i}^{\underline{\sigma}}$ among strategies in $\Sigma_{\tau_i}^k$, but among the larger set of strategies $\{\sigma_i \in \Delta(S_i) \mid \sigma_i(s_i) \geq \underline{\varepsilon}_{\tau_i, s_i}^k \ \forall s_i \in S_i\}$ as well. But then $\underline{\sigma}_{-\tau_i}^{k*}$ is a best response to $a_{\tau_i}^*$ in G for every $i \in N$ and $\tau_i \in \Upsilon_i$, establishing that $\underline{\sigma}^*$ is a sequential equilibrium of G . QED

Proof of Claim 9: Suppose that for some $h^B \in H/H^Z$ and $J \subset N$ there is a h^B -based nontrivial supported restriction B by J given R^* . Without loss of generality assume that $B_j \neq R_j^* \ \forall j \in J$.

If there is no $i \in N$ and $s_i, s'_i \in B_i$ such that s_i and s'_i reach h^B , but $s_i(h^B) \neq s'_i(h^B)$, then let h' be the unique information set such that $h^B = \text{imp}(h')$ and B reaches h' . If $h' \in H^Z$, then there is $i \in N$ such that B is a nontrivial supported restriction by $\{i\}$ given R^* , which by Claim 3 contradicts the definition of R^* . Therefore $h' \in H/H^Z$ and then trivially B is an h' -based supported restriction by J given R^* .

Assume now there is $i \in N$ and $s_i, s'_i \in B_i$ such that s_i and s'_i reach h^B , but $s_i(h^B) \neq s'_i(h^B)$. If $i \notin J$, then take any h' such that $h^B = \text{imp}(h')$. By the definition of a supported restriction B is an h' -based supported restriction by J given R^* . Assume next that $i \in J$. Let $H' = \{h' \in H \mid h^B = \text{imp}(h') \text{ and } \exists \theta_{-i} \in \Theta_{-i}^c(B_i) \text{ and } s_i \in BR_i(\theta_{-i}) \text{ such that } s_i \text{ reaches } h'\}$. For any $h' \in H'$ and $j \in N$ let $V_j(h') = \{s_j \in R_j^*/B_j \mid s_j \text{ reaches } h'\}$ and let $B'_j(h') = R_j^*/V_j(h')$. By construction $B'_j(h')$ is nonempty $\forall h' \in H'$ and $j \in N$. Let $B'(h') = \prod_{j \in N} B'_j(h')$ $\forall h' \in H'$. It cannot be that $B'(h') = R_j^* \ \forall h' \in H'$, since B is a nontrivial supported restriction given R^* .

Fix $h' \in H'$ and consider the restriction $B'(h')$ by J given R^* . For every $j \in J/\{i\}$ if $s_j \in BR_j(\theta_{-j})$, $\theta_{-j} \in \Theta_{-j}^c(B)$ and s_j reaches h^B , then also s_j reaches h' since j has a trivial action choice at h^B . Then for every $j \in J/\{i\}$ it holds that $B'_j(h') = B_j$.

Take any $\hat{s}_j \in BR_j(\theta_{-j})$, $\theta_{-j} \in \Theta_{-j}^c(B)$ such that \hat{s}_j reaches h^B . Since $B_j \neq R_j^*$, $\exists s_j \in R_j^*/B_j$ and $\theta'_{-j} \in \Theta_{-j}^c(R^*)$ such that $s_j \in BR_j(\theta_{-j})$. Then by the definition of a supported restriction $\exists \hat{\theta}_{-j} \in \Theta_{-j}^c(B)$ such that $\hat{s}_j \in BR_j(\hat{\theta}_{-j})$, and $\hat{\theta}_{-j}[h^B](s_{-j}) > 0$, s_{-j} reaching h^B implies that s_i reaches h' .

By the definition of a supported restriction $\hat{s}_j \in B_j$. Let now $s'_j \in R_j^*$ be such that $s'_j(h) = \hat{s}_j(h) \ \forall h \in h^B \cup \text{succ}(h^B)$. Since (R^1, \dots, R^*) is coherent, $\exists \theta'_{-j} \in \Theta_{-j}^c(R^*)$ such that $s'_j \in BR_j(\theta'_{-j})$. Let θ''_{-j} be such that $\theta''_{-j}[h] = \hat{\theta}_{-j}[h] \ \forall h \notin \text{succ}(h^B)$ and $\forall h \in h^B \cup h' \cup \text{succ}(h')$, and $\theta''_{-j}[h] = \theta'_{-j}[h] \ \forall h \in \text{succ}(h^B)/(h' \cup \text{succ}(h'))$. By construction $\theta''_{-j} \in \Theta_{-j}^c(R^*)$ and $s'_j \in BR_j(\theta''_{-j})$. Also by construction $u_j^h(s'_j, \theta''_{-j}) = u_j^h(\hat{s}_j, \hat{\theta}_{-j})$. The definition of a supported restriction then implies that $s'_j \in B_j$. This establishes that $B'(h')$ is an h' -based restriction by J given R^* . Since for every $j \in J/\{i\}$ it holds that $B'_j(h') = B_j$ and B is a supported restriction by J given R^* , $B'(h')$ is an h' -based supported restriction by J given R^* .

This concludes that if B is an h -based nontrivial supported restriction by J

given R^* , then $\exists \hat{h} \in \text{succ}(h)$ and C such that C is an \hat{h} -based nontrivial supported restriction by J given R^* . But since the game is of perfect information, that implies that there is $j \in J$ and C such that C is a nontrivial supported restriction by $\{j\}$ given R^* . By Claim 3 that contradicts the definition of R^* .

Therefore there is no nontrivial supported restriction given R^* . Claim 5 then implies that $O(R^*) = O(A^*)$. QED

Proof of Claim 10: For every $i \in N$, $\tau_i \in \Upsilon_i$, $s_i \in S_i$ and $\omega_{-i} \in \Omega_{-i}$ let $\hat{u}_{\tau_i}(s_i, \omega_{-i}) = \sum_{\underline{s}'_{-i} \in \underline{S}_{-i}} \hat{u}_{\tau_i}(s_i, \underline{s}'_{-i}) \omega_{-i}(\underline{s}'_{-i})$. Combining this with the definition of \hat{u}_{τ_i} and rearranging gives $\hat{u}_{\tau_i}(s_i, \omega_{-i}) = \sum_{v_{-i} \in \Upsilon_{-i}} [\sum_{\underline{s}; \underline{s}(v_{-i}) = t_{-i}} u_{\tau_i}(s_i, t_{-i}) \cdot \omega_{-i}(\underline{s}_{-i})] \cdot \varphi_i(\tau_i)(v_{-i}) = u_{\tau_i}(s_i, \omega_{-i})$. Note that this also implies that $s_i \in S_i$ is a best response to ω_{-i} for type τ_i of player i in G iff $s_i \in S_i$ is a best response to ω_{-i} for player τ_i in G' . But then the requirements for \underline{B} to be a supported restriction by J given \underline{A} are the same in G and G' . This implies that the intersection of all supported restrictions given any set $\underline{A} \subset \underline{S}$ such that $\underline{A} = \times_{i \in N} (\times_{\tau_i \in \Upsilon_i} \underline{A}_{\tau_i}) \neq \emptyset$ is the same in G and G' , implying $\underline{A}'^* = \underline{A}^*$. QED

11 References

- AMBRUS, A. (2001): "Coalitional rationality," *mimeo Princeton University*
- AUMANN, R. J. (1959): "Acceptable points in general cooperative n-person games," in *Contributions to the theory of games IV*, Princeton University Press, Princeton, NJ
- BATTIGALLI, P. (1996): "Strategic rationality orderings and the best rationalization principle," *Games and Economic Behavior*, 13, 178-200
- BATTIGALLI, P. (1997): "On rationalizability in extensive games," *Journal of Economic Theory*, 74, 40-61
- BATTIGALLI, P. and M. SINISCALCHI (2001): "Rationalization and incomplete information," *mimeo Bocconi University and Northwestern University*
- BATTIGALLI, P. and M. SINISCALCHI (2002): "Strong belief and forward induction reasoning," *Journal of Economic Theory*, 106, 356-391
- BENOIT, J. P. and V. KRISHNA (1985): "Finitely repeated games," *Econometrica*, 53, 905-922
- BENOIT, J. P. and V. KRISHNA (1993): "Renegotiation in finitely repeated games," *Econometrica* 61, 303-323
- BEN-PORATH, E. (1997): "Rationality, Nash equilibrium and backwards induction in perfect-information games," *Review of Economic Studies* 64, 23-46
- BEN-PORATH, E. and E. DEKEL (1992): "Signaling future actions and the potential for sacrifice," *Journal of Economic Theory*, 57, 36-51
- BERNHEIM, B. D. (1984): "Rationalizable strategic behavior," *Econometrica* 52, 1007-1028

- BERNHEIM, B. D., B. PELEG AND M. D. WHINSTON (1987): "Coalition-proof Nash equilibria I. Concepts," *Journal of Economic Theory*, 42, 1-12
- BERNHEIM, B. D. and D. RAY (1989): "Collective dynamic consistency in repeated games," *Games and Economic Behavior*, 1, 295-326
- BORGERS, T. (1991): "On the definition of rationalizability in extensive games," *mimeo University of California at San Diego*
- DEKEL, E. AND D. FUDENBERG (1990): "Rational behavior with payoff uncertainty," *Journal of Economic Theory*, 52, 243-267
- DUTTA, B. and R. VOHRA (2001): "Incomplete information, credibility and the core," *mimeo Brown University*
- FORGES, F. (1986): "An approach to communication equilibria," *Econometrica*, 54, 1375-1385
- FORGES, F., E. MINELLI and R. VOHRA (2002): "Incentives and the core of an exchange economy: a survey," *Journal of Mathematical Economics*, 38, 1-41
- FUDENBERG, D., D. KREPS and D. LEVINE (1988): "On the robustness of equilibrium refinements," *Journal of Economic Theory*, 44, 354-380
- GUL, F. and D. G. PEARCE (1996): "Forward induction and public randomization," *Journal of Economic Theory*, 70, 43-64
- HOLMSTROM, B. and R. B. MYERSON (1983): "Efficient and durable decision rules with incomplete information," *Econometrica*, 51, 1799-1819
- KAHN, C. M. and D. MOKHERJEE (1995): "Coalition proof equilibrium in an adverse selection insurance economy," *Journal of Economic Theory*, 66, 113-138
- KOHLBERG, E. and J.-F. MERTENS (1986): "On the strategic stability of equilibria," *Econometrica*, 54, 1003-1037
- KREPS, D. and R. WILSON (1982): "Sequential equilibria," *Econometrica*, 50, 863-894
- MAILATH, G., L. SAMUELSON and J. M. SWINKELS (1993): "Extensive form reasoning in normal form games," *Econometrica*, 61, 273-302
- MARX, L. M. and J. M. SWINKELS (1997): "Order independence and iterated weak dominance," *Games and Economic Behavior*, 18, 219-245
- MASKIN, E. and J. TIROLE (1990): "The principal-agent relationship with an informed principal: the case of private values," *Econometrica*, 58, 379-409
- MASKIN, E. and J. TIROLE (1992): "The principal-agent relationship with an informed principal, II: common values," *Econometrica*, 60, 1-42
- MYERSON, R. B. (1986): "Multistage games with communication," *Econometrica*, 54, 323-358
- MYERSON, R. B. (1989): "Credible negotiation statements and coherent plans," *Journal of Economic Theory*, 48, 264-303
- OSBORNE, M. J. (1990): "Signaling, forward induction and stability in finitely repeated games," *Journal of Economic Theory*, 50, 22-36
- OSBORNE, M. J. and A. RUBINSTEIN (1994): "A course in game theory," *MIT Press*
- PEARCE, D. G. (1984): "Rationalizable strategic behavior and the problem of perfection," *Econometrica*, 52, 1029-1050

- RENY, P. (1993): "Common belief and the theory of games with perfect information," *Journal of Economic Theory*, 59, 257-274
- RUBINSTEIN, A. (1980): "Strong perfect equilibrium in supergames," *International Journal of Game Theory*, 9, 1-12
- SELTEN, R. (1975): "Reexamination of the perfectness concept for equilibrium points in extensive games," *International Journal of Game Theory*, 4, 25-55
- SHIMOJI, M. (2002): "On forward induction in money-burning games," *Economic Theory*, 19, 637-648
- SHIMOJI, M. and J. WATSON (1998): "Conditional dominance, rationalizability and game forms," *Journal of Economic Theory*, 83, 161-195
- SOBEL, J., L. STOLE and I. ZAPATER (1990): "Fixed-equilibrium rationalizability in signaling games," *Journal of Economic Theory*, 52, 304-331
- VAN DAMME, E. (1989): "Stable equilibria and forward induction," *Journal of Economic Theory*, 48, 476-496
- VOHRA, R. (1999): "Incomplete information, incentive compatibility and the core," *Journal of Economic Theory*, 86, 123-147