Coalitional Rationalizability

Citation

Ambrus, Attila. 2006. Coalitional rationalizability. Quarterly Journal of Economics 121(3): 903-929.

Published Version

<http://dx.doi.org/10.1162/qjec.121.3.903>

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:3200266>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

COALITIONAL RATIONALIZABILITY*

ATTILA AMBRUS

This paper investigates how groups or coalitions of players can act in their collective interest in noncooperative normal form games even if equilibrium play is not assumed. The main idea is that each member of a coalition will confine play to a subset of their strategies if it is in their mutual interest to do so. An iterative procedure of restrictions is used to define a noncooperative solution concept, the set of coalitionally rationalizable strategies. The procedure is analogous to iterative deletion of never best response strategies, but operates on implicit agreements by different coalitions. The solution set is a nonempty subset of the rationalizable strategies.

I. INTRODUCTION

The main solution concept in noncooperative game theory, Nash equilibrium, requires stability only with respect to individual deviations by players. It does not take into account the possibility that groups of players might coordinate their moves, in order to achieve an outcome that is better for all of them. There have been several attempts to incorporate this consideration into the theory of noncooperative games, starting with the pioneering work of Schelling [1960] and Aumann [1959]. The latter offered strong Nash equilibrium, the first formal solution concept in noncooperative game theory that takes into account the interests of coalitions. More recently, Bernheim, Peleg, and Whinston [1987] proposed coalition-proof Nash equilibrium. This concept has since been used to derive predictions in a wide range of economic models. Examples include menu auctions [Bernheim and Whinston 1986], dynamic public good provision games [Chakravorti 1995], bankruptcy rules [Dagan, Serrano, and Volij 1997], principal-agent games [Gupta and Romano 1998], corporate takeovers [Noe 1998], common agency games [Konishi, LeBreton, and Weber 1999], and oligopoly competition [Delgado and Moreno 2004].

However, all of the concepts introduced so far are not being able to guarantee the existence of a solution in a natural class of

* I thank Dilip Abreu for support and extremely valuable comments at all stages of writing this paper. I would also like to thank Eddie Dekel, Drew Fudenberg, Faruk Gul, Wolfgang Pesendorfer, Ariel Rubinstein, and Marciano Siniscalchi for useful comments and discussions. Finally, I would like to thank Jeffrey Ely, Haluk Ergin, Erica Field, Jozsef Molnar, Wojciech Olszewski, Hugo Sonnenschein, Andrea Wilson, and seminar participants at various universities for useful comments.

games. This casts doubt on the validity of the solution they provide even in games in which a solution exists. Nonexistence is especially severe in the case of strong Nash equilibrium. Coalition-proof Nash equilibrium exists in a larger set of games, but at the cost of imposing debatable restrictions on which coalitions of players can make agreements with each other, and even so it cannot get around the nonexistence issue.

This paper proposes a new solution concept, coalitional rationalizability, to address the issue of coalitional agreements. Attention is restricted to normal-form games, although the principles proposed here can be applied to more general settings. Both a direct and a procedural definition for the solution set is provided. The latter is particularly easy to use. To obtain the set of coalitionally rationalizable strategies in any game, one can just use a simple iterative procedure.

The main conceptual innovation is that I depart from the usual equilibrium setting and present a nonequilibrium theory, like rationalizability [Bernheim 1984; Pearce 1984]. The coalitional agreements players can consider in this context take the form of restrictions of the strategy space. This means that players look for agreements to avoid certain strategies, without specifying play within the set of nonexcluded strategies. These agreements are more general than the ones that uniquely pin down a strategy profile for a coalition.

A restriction is supported if every group member always (for every possible expectation) expects a higher payoff if the agreement is made than if he instead chooses to play a strategy outside the agreement. If conjectures are associated with the payoffs that best response strategies to these conjectures yield, then the above requirement means that players in the group prefer any conjecture compatible with the agreement to any for which a strategy outside the agreement is a best response.

The construction assumes that players go through the following reasoning. First, every coalition of players looks for supported restrictions given the set of all strategies. Then players consider the set of strategy profiles that are consistent with all the above restrictions (their conjectures are restricted to be concentrated on these strategies) and look for further restrictions, given this smaller set of strategies. They continue this procedure and restrict the set of strategies further and further, until they reach a point at which there is no supported restriction by any coalition given the set of strategies that survived the procedure so far. This

procedure is analogous to iterative deletion of never best response strategies in that the order of restrictions does not matter. Furthermore, it deletes all strategies that the latter procedure eliminates. The new feature is that groups of players, as opposed to just individual players, can delete strategies.

The set of coalitionally rationalizable strategies is the set of profiles that survive the above procedure of iterated supported restrictions by coalitions.

The interpretation of the solution concept is that it is the set of outcomes which are compatible with the reasoning procedure reflected by the iterative procedure of supported restrictions. The latter restrictions come from introspection, based on the publicly known payoffs of the game. They reflect implicit agreements among players. In particular, they have to be self-enforcing: if a player believes that all others play according to a supported restriction, then it is strictly in her interest to play according to it.

The outline of the paper is as follows. Section II provides two simple examples to provide intuition on the logic of coalitional restrictions. The examples show that coalitional reasoning can lead to sharp predictions in certain games even when players' expectations are not assumed to be correct. In the first example it leads to an equilibrium profile (so the correctness of expectations is established as a result, as opposed to being assumed in the first place), while in the second one it significantly reduces the set of possible outcomes. Section III provides the formal construction of the theory, defining supported restrictions and the iterative procedure that obtains the solution set. It also contains the main results of the paper, establishing various properties of the iterative procedure and the solution set. Sections IV and V relate the solution set to other noncooperative solution concepts and examine connections with Pareto efficiency.

II. MOTIVATING EXAMPLES

This section provides examples on how coalitional reasoning can help players narrow down the set of possibilities they should consider, and to coordinate their action choices.

II.A. Voting with Costly Participation

Consider the following voting game. There are three potential voters and three possible outcomes. Assume that there is a status quo outcome S and that it can be changed only if at least

two of the voters show up and vote for the same new alternative. There are two potential new alternatives to vote for, *A* and *B*. Showing up and casting a vote costs $\varepsilon \in (0,1)$. Voters rank the possible outcomes the following way. Voter 1's favorite outcome is *A*, then *B*, and then *S*. Voter 2's favorite outcome is *A*, then *S*, and then *B*. Finally, voter 3's favorite outcome is *B*, then *S*, and then *A*. Assume that for every voter her favorite outcome yields a payoff of 2, her second favorite outcome yields 1, and her least favorite outcome has 0 (minus ε in each case if she showed up to vote). This game has multiple equilibria, and in fact for all three outcomes in the game there is some Nash equilibrium that yields that outcome. In one equilibrium all three voters stay at home, and the status quo outcome prevails. In another equilibrium voters 1 and 3 show up and vote for outcome *B*, while voter 2 stays home. And in yet another equilibrium voters 1 and 2 show up and vote for outcome *A*, while voter 3 stays at home. The above equilibria are not Pareto-ranked (note that each voter has a distinct least favorite outcome, so in any equilibrium someone's least favorite outcome is chosen). On the other hand, the following reasoning procedure selects a unique equilibrium in the game.

First, note that it is never rational (it is strictly dominated) for voter 3 to show up and vote for alternative *A*, since voting is costly and *A* is her least preferred outcome. Given that, the best possible outcome in the game for voters 1 and 2 is if they show up and vote for *A*. Although this point is fairly straightforward, let us make it more rigorous, to demonstrate the logic of supported restrictions that I formally define in the next section. Fix any conjecture concerning player 3's strategy that allocates zero probability to player 3 voting for *A*. Then the following are true. If players 1 and 2 both show up and vote for *A*, they expect a payoff of $2 - \varepsilon$ for sure. If player 1 shows up and votes for *B*, her expected payoff cannot be higher than $1 - \varepsilon$. If player 1 stays at home, her expected payoff cannot be higher than 1. The same are true for player 2, also (with the addition that voting for *B* is not even rational for her). Therefore, for any possible conjecture concerning player 3's strategy, both player 1 and player 2 are strictly better off coordinating on playing *A* than not coordinating and playing some other strategy. Completing the argument, if player 1 and player 2 act accordingly, then outcome *A* is indeed implemented, no matter whether player 3 stays at home or shows up

and votes for *B*. Given this, voter 3 should conclude that she is better off staying at home and not voting.

II.B. Dollar Division Game with External Reward

The second example demonstrates that even when coalitional reasoning does not lead to a unique prediction in a game, it can considerably narrow down the set of possible outcomes. Consider a classic dollar division game with the additional twist that players receive an external reward in the event that players behave “nicely.” Concretely, three players vote secretly and simultaneously on how to divide a dollar. If two or more players vote for the same allocation, the dollar is divided accordingly; otherwise the dollar is lost to the players. The added element is that if every player votes for allocations that would give all players in the game at least a quarter of a dollar, then every player gets an additional \$100 reward for the group being “generous,” independently of what happens to the original dollar (in particular, even in the event that it is not allocated to the players because of lack of agreement).¹ In this game coordinating on voting for allocations that give at least 1/4 dollar to every player is unambiguously mutually advantageous for the players. It is not clear how the original dollar should be divided, or whether it is reasonable to expect two or more players to vote for the same division, and if yes, then which players vote for the winning allocation. There is a conflict of interest among players regarding how to allocate the “last quarter,” but they have a strong incentive to propose at least 1/4 dollar to every player, since the external reward is much bigger than the stake that is to be divided. Players who coordinate along the line of common interest therefore should expect each other to propose allocations (x_1, x_2, x_3) such that $x_i \geq 1/4$, $\forall i \in \{1, 2, 3\}$, even if they are uncertain about exactly what allocations the others propose. It is worth pointing out that coalition-proof Nash equilibrium does not exist and therefore does not give any prediction in this game.²

1. This game can be interpreted as a simple model of the following situation. Several political parties in a postwar country trying to form a coalitional government, and an international organization makes a credible promise to provide a large amount of financial aid to the country if during the negotiations parties do not try to squeeze out any of the participants from power.

2. Section IV discusses the relationship between coalitional rationalizability and coalition-proof Nash equilibrium in detail, and briefly revisits this game.

III. CONSTRUCTION AND THE MAIN RESULTS

III.A. Notation and Basic Definitions

Let $G = (I, S, u)$ be a normal form game, where $I = \{1, \dots, n\}$ is the set of players, $S = \times_{i \in I} S_i$ is the set of strategies, and $u = \times_{i \in I} u_i$, $u_i: S \rightarrow R$, $\forall i \in I$ are the payoff functions. Assume that S_i is finite for every $i \in I$. Let $S_{-i} = \times_{j \in I \setminus \{i\}} S_j$, $\forall i \in I$, and let $S_{-J} = \times_{j \in I \setminus J} S_j$, $\forall J \subset I$. Similarly, for a generic $s \in S$, let $s_{-i} = \times_{j \in I \setminus \{i\}} s_j$, $\forall i \in I$, and let $s_{-J} = \times_{j \in I \setminus J} s_j$, $\forall J \subset I$. I will refer to non-empty groups of players (J such that $J \subset I$ and $J \neq \emptyset$) as *coalitions*.

I assume that players are Bayesian decision makers and that they can form correlated conjectures concerning other players' moves. Given the latter assumption, a strategy is a never best response if and only if it is strictly dominated (by a mixed strategy).³ Therefore, from this point on I use these terms interchangeably. Requiring conjectures to be independent (to be product probability distributions over the strategy space) does not change the qualitative results in the paper.

Let Δ_{-i} be the set of probability distributions over S_{-i} , representing the set of possible conjectures player i can have concerning other players' moves. For every $J \subset I$, $i \in J$ and $f_{-i} \in \Delta_{-i}$, let f_{-i}^{-J} be the marginal distribution of f_{-i} over S_{-J} .

The construction involves comparing expectations of players under different conjectures. For every $f_{-i} \in \Delta_{-i}$ and $s_i \in S_i$, let $u_i(s_i, f_{-i}) = \sum_{t_{-i}: t_{-i} \in S_{-i}} u_i(s_i, t_{-i}) \cdot f_{-i}(t_{-i})$ denote the expected payoff of player i if he has conjecture f_{-i} and plays pure strategy s_i .

Since players are Bayesian decision makers, the concept of best response plays a central role in what follows. For every $f_{-i} \in \Delta_{-i}$, let $BR_i(f_{-i}) = \{s_i | s_i \in S_i, u_i(s_i, f_{-i}) \geq u_i(t_i, f_{-i}), \forall t_i \in S_i\}$, the set of pure strategy best responses player i has against conjecture f_{-i} . For any $B \subset S$ such that $B \neq \emptyset$ and $B = \times_{i \in I} B_i$, let $\Omega_i^*(B_i) = \{f_{-i} | f_{-i} \in \Delta_{-i}, \exists b_i \in B_i \text{ such that } b_i \in BR_i(f_{-i})\}$. In words, $\Omega_i^*(B_i)$ is the set of beliefs that player i has against which there is a best response in B_i .

Let $\hat{u}_i(f_{-i}) = u_i(b_i, f_{-i})$ for any $b_i \in BR_i(f_{-i})$. Then $\hat{u}_i(f_{-i})$ is the expected payoff of a player if he has conjecture f_{-i} and plays

3. For the proof of this well-known result, see, for instance, Fudenberg and Tirole [1992, pp. 52–53].

a best response to his conjecture. That means $\hat{u}_i(f_{-i})$ is the expected payoff of a rational player if he has conjecture f_{-i} .

I will consider restrictions on the supports of players' conjectures. These restrictions are required to be product sets. For any A such that $A \subset S$ and $A = \times_{i \in I} A_i$, let $\Delta_{-i}(A) = \{f_{-i} | \text{supp} f_{-i} \subset A_{-i}\}$, the set of conjectures player i can have that are concentrated on A_{-i} (the set of conjectures that are consistent with player i believing that other players play inside A).

Certain product subsets of the strategy space play an important role in the construction below. These are sets that satisfy that whenever a player believes that others play inside the set then all her best responses are inside the set. Following standard terminology, I call these sets closed under rational behavior.⁴

DEFINITION. Set A is closed under rational behavior if it satisfies the following two properties:

$$(*) A = \times_{i \in I} A_i \quad \text{and} \quad A_i \subset S_i, \quad \forall i \in I$$

$$(**) BR_i(f_{-i}) \subset A_i, \quad \forall f_{-i} \in \Delta_{-i}(A), \quad \forall i \in I.$$

Let $\mathcal{M}$ denote the collection of sets closed under rational behavior.

I assume that players coordinate on restricting their play to a subset of the strategy space whenever by doing so each of them is guaranteed to get an expected payoff that is strictly higher than any expected payoff he could get if the restriction was not made and he played a strategy outside the restriction. These restrictions, called supported restrictions, constitute the main building block of the construction that follows.

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

DEFINITION. B is a supported restriction by J given A if it satisfies the following two properties:

$$(1) \quad B_i = A_i, \quad \forall i \notin J$$

$$(2) \quad \forall j \in J, f_{-j} \in \Omega_{-j}^*(A_j/B_j) \cap \Delta_{-j}(A)$$

it is the case that

$$\hat{u}_j(f_{-j}) < \hat{u}_j(g_{-j}) \quad \forall g_{-j} \in \Delta_{-j}(B) \quad \text{such that} \quad g_{-j}^{-j} = f_{-j}^{-j}.$$

4. The term was introduced by Basu and Weibull [1991].

Let $A \searrow_J B$ denote that B is a supported restriction by J given A . Let $A \searrow B$ denote $A \searrow_J B$ for some $J \subset I$, $J \neq \emptyset$.

Supported restrictions are defined given a nonempty product subset of the strategy space. The motivation is that players are certain that play is inside this set.⁵ Note that the set is allowed to be the set of all strategies. The first condition in the above definition requires that only the strategies of those players who are members of the given group be restricted. The second condition requires that for any player in the coalition, any belief to which he has a best response strategy outside the agreement yields a strictly lower expected payoff than any belief that is consistent with other players in the coalition keeping the agreement, holding the marginal expectation concerning the strategies of players outside the coalition fixed.

Intuitively, this definition considers the possible expected payoffs of a player if he chose to play a strategy to be excluded by the agreement, and compares them with payoffs he could expect if the restriction is made (if all the other players in the coalition confined their play to the restriction). Arguably, this is the most natural payoff comparison to use in deciding whether it is in the interest of a player to agree upon not playing the strategies to be excluded by a restriction. A restriction is then called supported if the expected payoffs that are compatible with the agreement strictly Pareto dominate those that are associated with playing strategies outside the restriction, for any fixed conjecture concerning players' strategies outside the coalition. In short, a restriction is supported if every coalition member prefers the agreement to playing a strategy outside the agreement, for every possible conjecture that he can have associated with the above two scenarios.

The reason marginal conjectures concerning the play of players outside the coalition are fixed is that since players make their moves secretly, the strategy choice of players outside the coalition cannot be made contingent on whether players in the coalition play inside B . The other players, after going through whatever mental procedure they use to formulate beliefs, may or may not believe that members of the coalition play according to the restriction. The point is that they do not have a chance to find this out. Note, however, that, as discussed above, the payoff compari-

5. For a rigorous formalization of this idea, see Ambrus [2005], where coalitional rationalizability is introduced in an epistemic framework.

son condition is required to hold for all possible conjectures concerning the play of players outside the coalition.

Besides its intuitive appeal, the definition above has the advantage that supported restrictions satisfy two desirable properties. One is that the concept is a generalization of eliminating never best response strategies by individual players (see Proposition 4 below), which ensures that a theory built on it is consistent with individual rationality. This property does not hold for definitions that are not based on comparing expected payoffs, for example if it is only required that for every player in the coalition every payoff within the restriction is strictly higher than every payoff outside the restriction. The second is that supported restrictions are self-enforcing. Note that the second condition in the above definition cannot hold if there is a player j in J and a conjecture f_{-j} which is concentrated on B_{-j} and against which j has a best response in A_{-j}/B_{-j} . This implies that if A is closed under rational behavior and A is a supported restriction given B , then B is also closed under rational behavior. If a player believes that the others play according to a supported restriction, then it is in his best interest to play according to the restriction too.

Nevertheless, comparing sets of payoffs is a far from obvious exercise, therefore one might want to consider alternative definitions to supported restriction. I do not take up that task here and instead refer the interested reader to Ambrus [2005].

III.B. The Set of Coalitionally Rationalizable Strategies and its Basic Properties

The construction below refers to sets that are closed under rational behavior and have the property that every strategy in the set is a best response to some conjecture concentrated on the set. I call these sets coherent.⁶

DEFINITION. Set A is coherent if it is closed under rational behavior and satisfies

$$(3) \quad \bigcup_{f_{-i} \in \Delta_{-i}(A)} BR_i(f_{-i}) = A_i, \quad \forall i \in I.$$

6. The terminology reflects that these sets satisfy the coherence requirement of Gul [1996] in that the strategies of a player that are inside the given set are implied by the restriction that other players play inside the set and that allowable beliefs about other players should include the convex hull of action profiles in the restriction. Investigating sets of strategies satisfying this type of requirement was first undertaken by Bernheim [1984] and Pearce [1984]. Basu and Weibull [1991] call these sets strictly closed under rational behavior.

Consider now the following iterative procedure of supported restrictions. Starting from the set of all strategies, in each step the intersection of all supported restrictions is retained. Below, it is shown that supported restrictions are compatible with each other in the sense that taking the intersection of all supported restrictions at any step of the procedure results in a nonempty set of strategies. This holds despite the fact that a player is part of many different coalitions and those coalitions may have different supported restrictions.

The procedure defined above can be thought of as a descriptive theory of belief formation. Players, based on the strategies and payoff functions of the game, look for supported restrictions given the set of all strategies. This means that at the beginning of the procedure they consider any strategy profile to be possible to be played. If such restrictions are found, then they expect the players in the corresponding coalitions to play inside the restrictions (to successfully coordinate their moves to play inside the restrictions). This requirement restricts the set of possible beliefs they can have. Then they look for supported restrictions with respect to the new, restricted set of possible beliefs. If such restrictions are found, then they expect players in the corresponding coalitions to play inside the restrictions, and so on, until the set of possible beliefs cannot be constrained any further. No other beliefs can be ruled out confidently based only on the information summarized in the payoff functions. I emphasize that, in the above interpretation, players do not explicitly make agreements with each other. They simply go through a reasoning procedure based on the commonly known payoff structure of the game and formulate their beliefs concerning the others' play according to this procedure. Explicit agreements would require preplay communication among players, which is not considered here.

For every $A \in \mathcal{M}$ let $\mathcal{F}(A)$ denote the collection of all supported restrictions given A .

DEFINITION. Let $A^0 = S$. For every $k \geq 1$ let $A_k = \bigcap_{B \in \mathcal{F}(A^{k-1})} B$.
Let $A^* = \bigcap_{k=0,1,2,\dots} A^k$.

The following two properties are useful in establishing one of the main results in the paper, nonemptiness of A^* . Proposition 1 establishes that the intersection of all supported restrictions given a set that is closed under rational behavior is nonempty. Proposition 2 establishes that if B is a supported restriction given

A and A is closed under rational behavior, then B remains a supported restriction given any set that is obtained from A by a sequence of supported restrictions. This property guarantees the internal consistency of the iterative procedure of supported restrictions, the main step in establishing that the order in which restrictions are made is inconsequential in this iterative procedure (see Proposition 5 below).

PROPOSITION 1. Let $A \in M$. Then $\bigcap_{B: B \in \mathcal{F}(A)} B \neq \emptyset$.

PROPOSITION 2. Let $A \in M$. Assume that $A \searrow_{J^B} B$. Let $C^0, \dots, C^k (k \geq 1)$ be such that $C^0 = A$ and $C^i \searrow_{J^i} C^{i-1}, \forall i = 1, \dots, k$. Then $C^k \searrow_{J^B} (B \cap C^k)$.

All formal proofs are in the Appendix. The important insight in establishing Proposition 1 is that for every player the strategies that are best responses for this player for her most optimistic conjecture (the one that yields the highest expected payoff) concentrated on A have to be included in every supported restriction by every coalition that this player is a member of. This follows from the definition of a supported restriction. For proving Proposition 2, the key step is showing that given a set that is closed under rational behavior, the intersection of two supported restrictions by the same coalition is a supported restriction itself, by the same coalition (and therefore, by Proposition 1, nonempty).

Proposition 3, a central result in this section, establishes that the iterative procedure of supported restrictions stops in a finite number of steps, and the set it obtains is nonempty, closed under rational behavior and has the property that, given this set, no coalition has a nontrivial supported restriction. The key step in proving this is showing, using Propositions 1 and 2, that A^k is nonempty and closed under rational behavior for every k . The finiteness of the game then implies all the claims in the proposition.

PROPOSITION 3. A^* satisfies the following properties:

- (i) nonempty
- (ii) closed under rational behavior
- (iii) $A^* \searrow B$ implies $B = A^*$
- (iv) $\exists K < \infty$ such that $A^k = A^*$ whenever $k \geq K$.

Although the iterative procedure is defined on pure strategies, allowing players to use mixed strategies would lead to the same pure strategies and some (not necessarily all) of the mixed

strategies with the same support. For more on this, see subsection VI.C.

To obtain some other properties of A^* , it is useful to identify supported restrictions by singleton coalitions. Proposition 4 establishes that these restrictions are equivalent to elimination of never best-response strategies.

PROPOSITION 4. Let $A \in M$ and $i \in I$. Then $A \searrow_{\{i\}} B$ iff $B = B_i \times A_{-i}$, $B_i \subset A_i$ and $s_i \in B_i/A_i$ implies that $\neg \exists f_{-i} \in \Delta_{-i}(A)$ such that $s_i \in BR_i(f_{-i})$.

Proposition 3 implies that, in particular, there is no supported restriction by any single-player coalition given A^* . Proposition 4 then implies that A^* satisfies (3). Furthermore, Proposition 3 establishes that $A^* \in \mathcal{M}$. Combining these results establishes that A^* is coherent. This immediately implies that A^* is contained in the set of rationalizable strategies, since the latter is the largest coherent set (see Bernheim [1984]).

A^* is defined to be the set obtained by a particular iterative procedure that requires supported restrictions to be made in a particular order, namely making all possible generalized supported restrictions simultaneously at every step. The next claim establishes that the particular order of restrictions does not matter. Any iterative procedure that makes some nontrivial supported restriction whenever one exists (for example, just making one restriction at a time, in any possible order) would yield the same solution set, A^* . This result is essentially a consequence of the one in Proposition 2.

PROPOSITION 5. Let $B^0 = S$. If there is no nontrivial supported restriction given B^0 , then let $B^1 = B^0$. Otherwise, let Θ^0 be a nonempty collection of nontrivial supported restrictions from B^0 , and let $B_1 = \bigcap_{B:B \in \Theta^0} B$. In a similar fashion once B^k is defined for some $k \geq 1$, let $B^{k+1} = B^k$ if there is no nontrivial supported restriction given B^k . Otherwise, let Θ^k be a nonempty collection of nontrivial supported restrictions given B^k , and let $B^{k+1} = \bigcap_{B:B \in \Theta^k} B$. Then there is $L \geq 0$ such that $B^k = A^*$, $\forall k \geq L$.

The last proposition of this section shows that the set of coalitionally rationalizable strategies is stable with respect to supported restrictions given any superset, and using this property gives a direct definition of the set.

Let $A = \times_{i \in I} A_i \neq \emptyset$ and $A \subset S$.

DEFINITION. A is externally coalitionally stable if for every $C \neq A$ such that $A \subset C$ it is the case that $A \subset \cap_{B: B \in \mathcal{F}(C)} B$ and $C \neq \cap_{B: B \in \mathcal{F}(C)} B$.

DEFINITION. A is internally coalitionally stable if $A \searrow B$ implies that $B = A$.

DEFINITION. A is coalitionally stable if it is both externally and internally coalitionally stable.

Intuitively, coalitional stability of A requires that whenever one starts out from a set larger than A , supported restrictions restrict that set "toward A ," while A itself cannot be restricted further.

PROPOSITION 6. A^* is the only coalitionally stable set in G .

The proof generalizes the arguments behind Proposition 2 and shows that (i) A^* is coalitionally stable; (ii) for any set A which is such that A/A^* is nonempty there is a superset of A given which there is a supported restriction that does not include A/A^* , implying that these sets cannot be coalitionally stable.

IV. RELATING COALITIONAL RATIONALIZABILITY TO OTHER SOLUTION CONCEPTS

I examine the connections between the set of coalitionally rationalizable strategies and some standard noncooperative solution concepts. These are Nash equilibrium, and the two most commonly used noncooperative coalitional equilibrium concepts: coalition-proof Nash equilibrium and strong Nash equilibrium.

IV.A. Nash Equilibrium

The set of coalitionally rationalizable strategies is not an equilibrium concept, so it is not surprising that it is not contained in the set of Nash-equilibrium outcomes (the outcomes that can be realizations of some mixed strategy Nash equilibrium). Consider the game of Figure I.

The only Nash equilibrium of the game is $(A2, B2)$. Nevertheless, the set of coalitionally rationalizable profiles is the whole game. For all other strategy profiles the sum of payoffs is negative. Still, if conjectures do not have to be correct, then other

	B1	B2	B3
A1	-2,1	-1,0	1,-2
A2	0,-1	0,0	0,-1
A3	1,-2	-1,0	-2,1

FIGURE I

strategies can be played since players can expect positive payoffs in the negative-sum matching pennies game $\{A1, A3\} \times \{B1, B3\}$.

Furthermore, as seen in previous examples, the set of Nash equilibria is not contained in the set of coalitionally rationalizable strategies. However, it is straightforward to show that there is always *at least one Nash equilibrium* of every game that is *inside the set of coalitionally rationalizable strategies*. This is a direct implication of the result that A^* is closed under rational behavior: any Nash equilibrium of the game with restricted strategy sets A^* (which as a finite normal-form game, has at least one Nash equilibrium in mixed strategies) is also a Nash equilibrium of G .

IV.B. Coalition-Proof Nash Equilibrium and Strong Nash Equilibrium

The first question addressed here is whether there can be nontrivial supported restrictions in games in which coalition-proof Nash equilibrium does not exist. The next example demonstrates that the answer is yes. Coalitional rationality can have bite in these games, in the sense that the set of coalitionally rationalizable strategies is strictly smaller than the set of rationalizable strategies.

The dollar division game is a classic example of a game in which coalitional equilibrium concepts do not exist. Consider the version of the game presented in Section II. The intuition that players should expect each other to vote for generous allocations so that the group can receive the external reward is captured by coalitional rationalizability. Proposing only allocations (x_1, x_2, x_3) such that $x_i \geq 1/4$, $\forall i \in \{1, 2, 3\}$ is a supported restriction for the coalition of all players given the set of all strategies, and it is the set of coalitionally rationalizable strategies. The game does not have any coalition-proof Nash equilibrium.

C1				C2					
		B1	B2	B3			B1	B2	B3
A1		2,1,0	0,0,0	-9,-9,-9	A1		1,2,0	0,2,1	-9,-9,-9
A2		2,0,1	1,0,2	-9,-9,-9	A2		0,0,0	0,1,2	-9,-9,-9
A3		-9,-9,-9	-9,-9,-9	-9,-9,-9	A3		-9,-9,-9	-9,-9,-9	-9,-9,-9

C3				
		B1	B2	B3
A1		-9,-9,-9	-9,-9,-9	-9,-9,-9
A2		-9,-9,-9	-9,-9,-9	-9,-9,-9
A3		-9,-9,-9	-9,-9,-9	-9,-9,-9

FIGURE II

The next question is whether the set of outcomes consistent with some coalition-proof Nash equilibrium is contained in the set of coalitionally rationalizable strategies. The following example demonstrates that the answer is no.

In the game of Figure II, $(A3, B3, C3)$ is the unique coalition-proof equilibrium (even allowing for mixed strategies). It is straightforward to establish that there is no self-enforcing profile in which players only play strategies inside $\{A1, A2\} \times \{B1, B2\} \times \{C1, C2\}$ with positive probability, by showing that from every such Nash equilibrium there is a two-player coalition that can profitably deviate to another Nash equilibrium.⁷ Next, observe that in any Nash equilibrium a player can only play his third strategy with positive probability if at least one of the other players plays his third strategy with probability 1 (otherwise, the third strategy cannot be a best response). But then in any self-

7. In a three-player game a profile is self-enforcing iff it satisfies the following two properties: (i) it is a Nash equilibrium; (ii) no coalition of two players can deviate to a profile that is a Nash equilibrium in the component game induced on them by the third player's strategy in such a way that both of them are strictly better off. For the general definition of self-enforcing, see Bernheim, Peleg, and Whinston [1987].

enforcing profile at least two players have to play their third strategies with probability 1. It is straightforward to show that there is no self-enforcing profile in which two players play their third strategies with probability 1 and the third player does not, because then the first two players have a joint deviation from which neither of them could deviate further profitably. Hence, the only candidate even in mixed strategies for a self-enforcing profile is when player 1 plays $A3$ with probability 1, player 2 plays $B3$ with probability 1, and player 3 plays $C3$ with probability 1. Furthermore, it is a self-enforcing profile since no single-player or two-player coalition can have a profitable deviation, and the coalition of all three players does not have a self-enforcing deviation, because there is no other self-enforcing profile in the game. Since $(A3, B3, C3)$ is the only self-enforcing profile in the game, it is the unique coalition-proof Nash equilibrium.

Furthermore, the set of coalitionally rationalizable strategies is $\{A1, A2\} \times \{B1, B2\} \times \{C1, C2\}$ (this is the set obtained after the first round of the iterative procedure at which point there is no more nontrivial supported restriction), so $(A3, B3, C3)$ is not coalitionally rationalizable. In fact, the set of coalitionally rationalizable strategies and the set of coalition-proof equilibria are disjoint sets in this game.

Note that $(A3, B3, C3)$ is a coalition-proof equilibrium only because the game with the set of strategies $\{A1, A2\} \times \{B1, B2\} \times \{C1, C2\}$ does not have a coalition-proof equilibrium. On the other hand, all players would strictly prefer to switch play to that part of the game no matter what happens over there (they cannot agree upon a concrete profile, but they would all agree to restricting their moves to $\{A1, A2\} \times \{B1, B2\} \times \{C1, C2\}$). This prediction is clearly more reasonable than that players will play the profile $(A3, B3, C3)$.

The above game could easily be made generic by some small perturbation of the payoffs, such that the set of coalitionally rationalizable strategies and the set of coalition-proof equilibria remain the same. Thus, even in generic games a coalition-proof equilibrium might not be coalitionally rationalizable.

The next proposition shows that only the nonexistence of coalition-proof Nash equilibrium in some restriction of the original game can result in some coalition-proof Nash equilibrium not being coalitionally rationalizable. Proposition 7 establishes that if in every restriction of the original game there exists a coalition-proof Nash equilibrium, then all the pure strategy coalition-proof

Nash equilibria of the original game are coalitionally rationalizable. The main argument in the proof is as follows. If a profile is not included in the set of coalitionally rationalizable strategies, then there is k and a coalition J such that the profile is included in A^k , but there is a supported restriction B by J given A such that the profile is not included in B . Then it can be shown that coalition J could deviate in a credible and profitable manner from the original profile to any coalition-proof Nash equilibrium of the game restricted to B , contradicting that the profile is a coalition-proof Nash equilibrium itself. The result can be extended to mixed strategy coalition-proof Nash equilibria by requiring that every game which is obtained from the original game by fixing a mixed strategy profile for some players and restricting the strategies of the other players to a subset of their original strategy set has a coalition-proof Nash equilibrium.

For every $A \subset S$ such that $A = \times_{i \in I} A_i \neq \emptyset$, let $G[A] = (I, A, u_A)$ denote the normal-form game for which $u_A(s) = u(s)$, $\forall s \in A$. It is the game obtained from G by restricting the set of strategies to A .

PROPOSITION 7. If for every $A \subset S$ such that $A = \times_{i \in I} A_i \neq \emptyset$, it holds that $G[A]$ has a coalition-proof Nash equilibrium, then every pure strategy coalition-proof Nash equilibrium of G is contained in A^* .

Proposition 7 and the example of Figure II illustrate that the issue of nonexistence confounds the iteratively defined solution concept of coalition-proof Nash equilibrium. This is not surprising given the iterative definition of the concept. The set of coalition-proof Nash equilibria of a game directly depends on coalition-proof Nash equilibria of restrictions of the original game. Therefore, coalition-proof Nash equilibrium can only be trusted to give a reasonable prediction if it gives a reasonable prediction in every restriction of the game (namely if it exists in every restriction). It is unclear to me whether in the latter class of games coalition-proof Nash equilibria give more reasonable predictions than other Nash equilibria that lie within the set of coalitionally rationalizable strategies. This is an issue I hope to return to in future work.

The section concludes by investigating relations to strong Nash equilibria: profiles that satisfy that no coalition has any profitable joint deviation. The latter is a very strong stability requirement, resulting in nonexistence of a strong Nash equilib-

rium in many games. The next proposition establishes that, unlike coalition-proof Nash equilibria, every strong Nash equilibrium of every game must be contained in the set of coalitionally rationalizable strategies. This means that there is no inherent contradiction between supported restrictions in a nonequilibrium setting and group deviations in an equilibrium setting. Inconsistencies only arise if the set of allowable coalitional deviations are restricted according to the definition of coalition-proof Nash equilibrium.

PROPOSITION 8. Let $\sigma = (\sigma_1, \dots, \sigma_I)$ be a strong Nash equilibrium profile. Then $\text{supp}\sigma \subset A^*$.

V. PARETO EFFICIENCY

Coalitional rationalizability takes the interest of coalitions other than the coalition of all players into account. Therefore, in general it does not guarantee Pareto efficiency. There is no containment relationship between the set of coalitionally rationalizable strategies and the set of Pareto undominated profiles, or rationalizable profiles that are Pareto undominated by other rationalizable profiles, or Nash equilibria that are Pareto undominated by other Nash equilibria.

The fact that groups of players expect each other to pursue common gains can make all of them worse off, as the game of Figure III shows.

In this game the (strict) Nash equilibrium profile $(A1, B1, C1)$ is not coalitionally rationalizable because no matter what player 3 does, playing $(A2, B2)$ always gives the highest payoff for players 1 and 2. But then player 3 is better off playing $C2$, making $(A2, B2, C2)$ the only coalitionally rationalizable profile, which is strictly Pareto-dominated by $(A1, B1, C1)$.

This example might seem puzzling, given that coalitional rationalizability is built on the assumption that players try to attain common gains. However, if not only the coalition of all players, but also subcoalitions of players act along these lines, then nothing guarantees Pareto efficiency of the resulting outcome. The fact that coalitions cannot commit not to go for a common gain can make them worse off. I view this as a coalitional version of the insight that is obtained from the prisoner's dilemma, where the fact that players cannot commit not to play individually rational strategies makes both of them worse off.

	C1			C2	
	B1	B2		B1	B2
A1	2,2,2	0,0,0	A1	0,0,0	0,0,0
A2	0,0,0	3,3,0	A2	0,0,0	1,1,1

FIGURE III

Although in the above example Pareto inefficiency results from the fact that there is a highlighted coalition that cannot commit not to go for coalitional gains, it can be shown that inefficiency can arise in perfectly symmetric games as well.⁸

Pareto efficiency can be guaranteed only in special classes of games. For example, it is easy to establish that in games that have a Pareto dominant profile among rationalizable strategies, that profile is the unique element in the set of coalitionally rationalizable strategies. Also, in two-player games the support of every Pareto-undominated Nash equilibrium is contained in the set of coalitionally rationalizable strategies. The proofs of these claims are straightforward and therefore omitted.

VI. DISCUSSION

VI.A. *Epistemic Definition*

The definition of coalitional rationalizability provided in this paper is along the lines of the original definition of rationalizability (see Pearce [1984] and Bernheim [1984]). It is characterized by an iterative procedure which has an intuitive interpretation: players expect all supported restrictions to be made given the set of all strategies, and they expect all supported restrictions to be made given the set of strategies that are consistent with the above requirement, and so on. The main advantage of this definition, besides that it refers to an intuitive reasoning procedure, is that it is constructive, making the set of rationalizable strategies relatively easy to compute in examples and applications. However, it does not directly reveal what primitive assumptions

8. See the previous version of the paper for an example.

on players' beliefs and action choices imply that they play coalitionally rationalizable strategies. The set of rationalizable strategies was subsequently shown to be equivalent to the set of strategies compatible with rationality and common certainty of rationality [Tan and Werlang 1988; Brandenburger and Dekel 1993], using the formalism of interactive epistemology. This raises the question whether the set of coalitionally rationalizable strategies has a similar interpretation. The direct characterization of Proposition 6 is a step in this direction. I do not pursue the issue further in this paper, instead referring the interested reader to Ambrus [2005]. That paper provides alternative, epistemic definitions of coalitional rationalizability.

VI.B. Mixed Strategies

Coalitional rationalizability is a concept defined on pure strategies (with respect to the players' actions, a player's conjecture can be any probability distribution on the other players' strategy set). However, the construction is also valid if players are allowed to play mixed strategies. It is straightforward to extend the concept of supported restriction and then coalitional rationalizability to the space of mixed strategies. One can then show that coalitionally rationalizable mixed strategies are a subset of mixed strategies whose support is inside the set of coalitionally rationalizable pure strategies.⁹ Furthermore, they include all coalitionally rationalizable pure strategies.

VI.C. Preplay Communication

A natural issue to investigate is the interaction of preplay communication and coalitional rationality. The particular questions that arise include whether all supported restrictions remain credible in a context with communication among players, whether there are new restrictions that become credible and whether the set of solutions compatible with the theory still has product structure. Since these issues are complicated and possibly depend on the exact specification of the rules of communication, I leave the task of formally addressing the problem of coalitional agreements in a framework with communication to a future project.

Here I just note that there are reasons to believe that preplay

9. The inclusion can be strict. This is analogous to the relationship between rationalizable pure and mixed strategies.

communication can have a role in determining whether or not players use coalitional reasoning. Coalitional rationalizability requires confidence that other players reason a particular way. Preplay communication can help establish the necessary amount of trust for the restrictions involved. Experimental game theory provides some support for this claim. In certain coordination games, preplay communication increases players' propensity to play Pareto optimal outcomes, and multisided preplay communication may increase cooperation more than one-sided communication (see Cooper et al. [1992] and Charness [2000]).

VII. RELATED LITERATURE

Section IV related coalitional rationalizability to various equilibrium concepts. Besides the general solution concepts mentioned, there are several contributions in the literature that incorporate coalitional reasoning into the play of normal-form games in more specific settings.

Chwe [1994], Mariotti [1997], and Xue [1998, 2000] assume that players engage in a possibly infinite negotiation procedure before playing a normal-form game. These models are similar to the construction behind coalitional rationalizability in that coalitions can freely form and that binding agreements are not available. The main difference, besides that only point agreements are considered, is the assumption that coalitions act publicly and therefore agreements are publicly observed.

Noncooperative coalitional bargaining considers extensive form noncooperative games to model n -player coalitional bargaining situations based on characteristic function games (or generalizations of those). These characteristic forms can be derived from normal-form games, as done in Ray and Vohra [1997, 1999]. The main difference between this approach and the one presented in this paper is the central assumption in the above papers that members inside a coalition can make binding agreements, and only the play among different coalitions occurs in a noncooperative fashion. The current paper conforms with the tradition of noncooperative game theory and retains the assumption that players cannot make any binding agreements.

Rabin's concept of Consistent Behavioral Theories (see Rabin [1994b]) can be used to incorporate coalitional reasoning into normal-form games, and the paper proposes one such theory, Pareto-focal rationalizability. However, while Rabin's approach

starts with some exogenously given set of focal points, the procedure in this paper can endogenously explain why certain outcomes are focal in a game.

There are papers investigating the role of preplay communication before playing a normal-form game, examining whether it leads to the type of belief restrictions considered in this paper. Farrell [1988] assumes that before playing a normal-form game one player can send a suggestion to the others. This suggestion is allowed to be a set of strategies, not just a single profile. He points out that there are games in which players clearly do not want to make any single-profile agreement, although his considerations are different than the ones highlighted in this paper. Rabin [1990] considers assumptions similar to those of Farrell in a rationalizability setting. Watson [1991] introduces a model in which one player can suggest playing inside some subset of the strategy space. The definition of when this message is credible is somewhat similar to the definition of supported restriction, although it is not a generalization of best response strategies. Furthermore, Watson's concept is defined only for two-player games, in which the issues of coalitional agreements are fairly simple.

VIII. CONCLUSION

In a lot of economic and political situations, subgroups of the participants have an incentive to coordinate their action choices. Various coalitional equilibrium concepts were proposed in the literature that assumed that players with similar interest could coordinate their play. However, the concepts proposed so far cannot guarantee existence in a natural class of games. The way I interpret this is that allowing coalitions to be able to make only point restrictions (agreements) leads to logical inconsistencies. As this paper shows, these inconsistencies do not have to arise if one allows for set-valued restrictions, suggesting that the latter is the appropriate framework to incorporate coalitional reasoning into noncooperative game theory.

APPENDIX

Proof of Proposition 1. Let a be such that $u_j(a) = \max_{s:s \in A} u_j(s)$. Then by the definition of a supported restriction, $a_j \in B_j \forall B$ such

that $A \searrow B$ because a_j is a best response against a_{-j} (it yields the maximum payoff in A and $A \in \mathcal{M}$). Therefore, $\bigcap_{B: B \in \mathcal{T}(A)} B_j \neq \emptyset$. This establishes the claim since j was arbitrary and $\bigcap_{B: B \in \mathcal{T}(A)} B$ is a product set. QED

LEMMA 1. Let $A \in M$ and $A \searrow_J B$. Then $B \in M$.

Proof of Lemma 1. Suppose that $B \notin \mathcal{M}$. Then $\exists j, a_j$ and f_{-j} such that $j \in J$, $a_j \in A_j/B_j$, $f_{-j} \in \Delta_{-j}(B)$ and $a_j \in BR_j(f_{-j})$, which contradicts $A \searrow_J B$. QED

LEMMA 2. Let $A \searrow_J B$, and let C be such that $C \in M$, $C \subset A$ and $C \cap B \neq \emptyset$. Then $C \searrow_J (C \cap B)$.

Proof of Lemma 2. Lemma 1 and $C \in \mathcal{M}$ imply that $B \in \mathcal{M}$. Then $C \cap B \in \mathcal{M}$ by the definition of a set closed under rational behavior. Let j and c_j be such that $j \in J$ and $c_j \in C_j/B_j$. Note that $A \searrow_J B$ implies that $u_j(c_j, f_{-j}) < u_j(b_j, g_{-j})$, $\forall b_j, f_{-j}, g_{-j}$ such that $f_{-j} \in \Omega_{-j}^*(A/B) \cap \Delta_{-j}(A)$, $c_j \in BR_j(f_{-j})$, $g_{-j} \in \Delta_{-j}(B)$, $b_j \in BR_j(g_{-j})$ and $g_{-j}^{-J} = f_{-j}^{-J}$, $\forall s_{-j} \in S_{-j}$. But since $C \subset A$ and $C \cap B \subset B$, this implies that $u_j(c_j, f_{-j}) < u_j(b_j, g_{-j})$, $\forall b_j, f_{-j}, g_{-j}$ such that $f_{-j} \in \Omega_{-j}^*((C \cap B)) \cap \Delta_{-j}(C)$, $c_j \in BR_j(f_{-j})$, $g_{-j} \in \Delta_{-j}(C \cap B)$, $b_j \in BR_j(g_{-j})$ and $g_{-j}^{-J} = f_{-j}^{-J}$, $\forall s_{-j} \in S_{-j}$. Since this holds for every j and c_j such that $j \in J$ and $c_j \in C_j/B_j$, $(C) \searrow_J (C \cap B)$. QED

Proof of Proposition 2. By Lemma 1, $B \in \mathcal{M}$, and $C^i \in \mathcal{M}$ $\forall i = 1, \dots, k$. By Proposition 1, $B \cap C^1 \neq \emptyset$. Then by Lemma 2, $C^1 \searrow_{J^B} (B \cap C^1)$. Now suppose that $C^n \searrow_{J^B} (B \cap C^n)$ for some $1 \leq n \leq k-1$. Then by definition $B \cap C^n \neq \emptyset$. Then Proposition 1, $C^n \searrow_{J^B} (B \cap C^n)$, and $C^n \searrow_{J^n} C^{n-1}$ imply that $B \cap C^{n+1} \neq \emptyset$. Then Lemma 2 implies that $C^{n+1} \searrow_{J^B} (B \cap C^{n+1})$. The claim follows by induction. QED

Proof of Proposition 3. Since S is finite and $A^{k-1} \supset A^k$, $\forall k \geq 1$, the second part of the claim is immediate. Note that $A^0 = S \in \mathcal{M}$. Now assume that $A^k \in \mathcal{M}$ for some $k \geq 0$. By Proposition 1, $A^{k+1} \neq \emptyset$. By Lemma 2, A^{k+1} can be reached from A^k by a sequence of supported restrictions, and then by Proposition 2, $A^{k+1} \in \mathcal{M}$. By induction, $A^k \neq \emptyset$, and $A^k \in \mathcal{M}$, $\forall k \geq 0$. Since $A^* = A^k$ whenever $k \geq K$, this implies that $A^* \neq \emptyset$ and $A^* \in \mathcal{M}$. Now suppose that there exists a nontrivial supported restriction given A^* . Since $A^* = A^K$, this implies that there is a

nontrivial supported restriction given A^K , which contradicts that $A^{K+1} = A^K$. QED

Proof of Proposition 4. This proof follows from the fact that for a single-player coalition $\{i\}$, requirement 2 in the definition of supported restriction is equivalent to requiring that there are no s_i and f_{-i} such that $s_i \in B_i/A_i$, $f_{-i} \in \Delta_{-i}(A)$ and $s_i \in BR_i(f_{-i})$. QED

Proof of Proposition 5. Since the sequence of sets $(B^k)_{k=0}^\infty$ is nested and S is finite, there is $L \geq 0$ such that $B^k = B^L$, $\forall k \geq L$. Since $B^0 = S$, $B^0 \in \mathcal{M}$. Now assume that $B^k \in \mathcal{M}$ for some $k \geq 0$. By Proposition 1, $\bigcap_{B: B \in \mathcal{T}(B^k)} B \neq \emptyset$, so $B^{k+1} = \bigcap_{B: B \in \mathcal{T}(B^k)} B \neq \emptyset$. By Lemmas 1 and 2, $B^{k+1} \in \mathcal{M}$. Since $B^k = B^L$, $\forall k \geq L$, by definition of the sequence $(B^k)_{k=0}^\infty$ there is no nontrivial supported restriction given B^L .

By definition $B^L \subset A^0$. Note that by Lemma 2 and Proposition 2 B^L can be reached from A^0 by a sequence of supported restrictions. Then by Lemma 1 $B^L \in \mathcal{M}$. Therefore, by Proposition 2 $A^0 \searrow B$ implies that $B^L \searrow B^L \cap B$. Then since there is no nontrivial supported restriction given B^L , $B^L \subset A^1$. Then by Lemma 2 B^L can be reached from A^1 by a sequence of supported restrictions. An inductive argument shows that $B^L \subset A^k$, $\forall k \geq 0$, and therefore $B^L \subset A^*$. A symmetric argument establishes that $A^* \subset B^k$, $\forall k \geq 0$, and therefore $A^* \subset B^L$. QED

LEMMA 3. Let A be such that $A^* \subset A$ and $A \neq A^*$. Then $A^* \subset \bigcap_{B: B \in \mathcal{T}(A)} B$.

Proof of Lemma 3. Suppose not. Then there are $B \subset A$ and $J \subset I$ such that $A \searrow_J B$ and $A^* \not\subseteq B$. First, consider $B \cap A^* \neq \emptyset$. Then by Lemma 2, $A^* \searrow_J (B \cap A^*)$, contradicting that there is no nontrivial supported restriction given A^* .

Now consider $B \cap A^* = \emptyset$. Then there is $k \geq 0$ such that $B \cap A^k \neq \emptyset$ and $B \cap A^{k+1} = \emptyset$. As established above, $A^k \in \mathcal{M}$. Together with $A \in \mathcal{M}$ this implies that $A \cap A^k \in \mathcal{M}$, because for any $i \in I$ and any $f_{-i} \in \Delta_{-i}(A \cap A^k)$, $BR_i(f_{-i}) \in A$ since $A \in \mathcal{M}$ and $BR_i(f_{-i}) \in A^k$ since $A^k \in \mathcal{M}$, so $BR_i(f_{-i}) \in A \cap A^k$. Since $A \searrow_J B$, by Proposition 2, $(A \cap A^k) \searrow_J (B \cap A^k)$. Furthermore, since $A \cap A^{k+1} \neq \emptyset$ (they both contain A^*), $(A \cap A^k) \searrow_J (A \cap C) \forall C \in \mathcal{T}(A^k)$. The above implies that $\bigcap_{C: C \in \mathcal{T}(A \cap A^k)} C \subset (A \cap A^{k+1}) \cap (B \cap A^k)$. But $(B \cap A^k) \cap (A \cap A^{k+1}) = B \cap A^{k+1} = \emptyset$, contradicting Proposition 1. QED

LEMMA 4. Let $A \in M$ be such that $A/A^* \neq \emptyset$ and $A \cap A^* \neq \emptyset$.
Then $\bigcap_{B: B \in \mathcal{F}(A)} B \neq A$.

Proof of Lemma 4. There exists $k \geq 0$ such that $A \subset A^k$ and $A \not\subseteq A^{k+1}$. Then there are $B \subset A$ and $J \subset I$ such that $A^k \searrow_J B$ and $A \cap B \neq A$. Note that $A^k \searrow_J B$ implies $A^{k+1} \subset B$ which in turn implies $A^* \subset B$. Therefore, $A \cap A^* \neq \emptyset$ implies that $A \cap B \neq \emptyset$. Furthermore, $A^k \searrow_J B$ implies that $A \in \mathcal{M}$, which together with $A \in \mathcal{M}$ and $A \cap B \neq \emptyset$ implies that $A \cap B \in \mathcal{M}$. Then by Lemma 2 $A \searrow_J (A \cap B)$, so $\bigcap_{B: B \in \mathcal{F}(A)} B \neq A$. QED

Proof of Proposition 6. Lemmas 3 and 4 establish that A^* is externally coalitionally stable. Proposition 3 establishes that A^* is internally coalitionally stable. By Lemma 4, if A is such that $A/A^* \neq \emptyset$ and $A \cap A^* \neq \emptyset$, then A cannot be coalitionally stable. If $A/A^* = \emptyset$, then $A \subset A^*$, in which case either $A = A^*$ or A cannot be coalitionally stable, since A^* contains it and there is no nontrivial supported restriction given A^* . And if $A \cap A^* = \emptyset$, then there is $k \geq 0$ such that $A \subset A^k$ and $A \not\subseteq A^{k+1}$. Therefore, A is not coalitionally stable, since $A \subset A^k$ and $A \not\subseteq \bigcap_{B: B \in \mathcal{F}(A^k)} B$. QED

Proof of Proposition 7. Suppose that s^* is a pure strategy coalition-proof Nash equilibrium and $s^* \notin A^*$. Then there is $k \geq 0$ such that $s^* \in A^k$ but $s^* \notin A^{k+1}$. That implies that there is $J \subset I$ and $B \subset A$ such that $A^k \searrow_J B$ and $s^* \notin B$. Let s' be a coalition-proof Nash equilibrium of $G[B_{J'} \times s_{I/J'}^*]$. The starting assumption guarantees the existence of such profile s' . Note that for every $j \in J$ it holds that s'_j is a best response against a conjecture that allocates probability 1 to s'_{-j} , and that s_j^* is a best response against a conjecture that allocates probability 1 to s_{-j}^* . This implies that $u_j(s') > u_j(s^*)$, $\forall j \in J'$ because B is a supported restriction by J given A^k . Therefore, s' is a profitable coalitional deviation from s^* by J' . Suppose now that there is a credible profitable coalitional deviation s'' from s' by $J'' \subset J'$. The credibility of this deviation implies that for every $j \in J''$, s''_j is a best response against the belief that allocates probability 1 to s''_{-j} . Consider first the case that $s''_j \in A_j^k$, $\forall j \in J''$. Since s' is a coalition-proof Nash equilibrium of $G[B_{J'} \times s_{I/J'}^*]$, it can only be that $s''_i \notin B_i$ for some $i \in J''$. But then the profitability of this deviation implies that $u_i(s'') > u_i(s')$, contradicting $A^k \searrow_J B$, since $i \in J'' \subset J$. Next consider the case that $s''_j \notin A_j^k$ for some $j \in J''$. Let $m \geq 0$ be such that $s''_j \in A_j^m$, $\forall j \in J''$, but there is $i \in J''$ such that $s''_i \notin A_i^{m+1}$. Note that $m < k$ since $s''_j \notin A_j^k$. Then there

is some $L \subset I$ and $C \subset A^m$ such that $A^m \searrow_L C$ and $s''_i \notin C_i$. But the profitability of the above deviation implies that $u_i(s'') > u_i(s')$, contradicting $A^m \searrow_L C$. This establishes that s' is a credible profitable deviation from s^* by J' since there is no credible profitable deviation from it by a subcoalition of J' . This contradicts that s^* is a pure strategy coalition-proof Nash equilibrium. QED

Proof of Proposition 8. Let $A_i = \text{supp}\sigma_i$, $\forall i \in I$, and let $A = \times_{i \in I} A_i$. Suppose that $A \not\subseteq A^*$. Then there is $k \geq 0$ such that $A \subset A^k$, but $A \not\subseteq A^{k+1}$. This implies that there are $B \subset A^k$ and $J \subset I$ such that B is a supported restriction by J given A^k , and $A \not\subseteq B$. Let $L = \{j | j \in J, \exists s_j \text{ such that } s_j \in A_j \text{ and } s_j \notin B_j\}$. For every $l \in L$, let a_l be such that $a_l \in A_l$ and $a_l \notin B_l$. For every $l \in L$, let f_{-l} be the conjecture of player l corresponding to the others playing the profile $\sigma_{-l}: f_{-l}(s_{-l}) = \times_{i \in I \setminus \{l\}} \sigma_i(s_i)$. Note that $a_l \in BR_l(f_{-l})$, $\forall l \in L$. Now let G_L be the truncated game in which the set of players are L , the set of strategies are B_l , $l \in L$, and the payoff functions are $\hat{g}_l(s_L) = g_l(s_L, \sigma_{-L})$. Since its strategy sets are compact and payoff functions are continuous, G_L has a Nash equilibrium in mixed strategies. Let $\hat{\xi}_L$ be such a profile. Since $B \in \mathcal{M}$ by Lemma 1, for every $l \in L$, $\hat{\xi}_l$ is a best response against the profile $(\hat{\xi}_{L \setminus l}, \sigma_{-L})$. Then since B is a supported restriction by J given C , $u_l(\hat{\xi}_L, \sigma_{-L}) > u_l(a_l, \sigma_{-L}) = u_l(\sigma)$. But that implies $\hat{\xi}_L$ is a profitable deviation for L from σ , contradicting that σ is a strong Nash equilibrium. QED

PRINCETON UNIVERSITY AND HARVARD UNIVERSITY

REFERENCES

- Ambrus, Attila, "Theories of Coalitional Rationality," mimeo, Harvard University, 2005.
- Aumann, Robert J., "Acceptable Points in General Cooperative n-person Games," in Albert W. Tucker and R. Duncan Luce, eds., *Contributions to the Theory of Games IV* (Princeton, NJ: Princeton University Press, 1959), pp. 287–324.
- Basu, Kaushik, and Jorgen W. Weibull, "Strategy Subsets Closed under Rational Behavior," *Economics Letters*, XXXVI (1991), 141–146.
- Bernheim, B. Douglas, "Rationalizable Strategic Behavior," *Econometrica*, LII (1984), 1007–1028.
- Bernheim, B. Douglas, Bezalel Peleg, and Michael D. Whinston, "Coalition-Proof Nash Equilibria I. Concepts," *Journal of Economic Theory*, XVII (1987), 1–12.
- Bernheim, B. Douglas, and Michael D. Whinston, "Menu Auctions, Resource Allocation, and Economic Influence," *Quarterly Journal of Economics*, CI (1986), 1–31.
- Brandenburger, Adam, and Eddie Dekel, "Hierarchies of Beliefs and Common Knowledge," *Journal of Economic Theory*, LIX (1993), 189–198.

- Chakravorti, Bhaskar, "Dynamic Public Goods Provision with Coalitional Manipulation," *Journal of Public Economics*, LVI (1995), 143–161.
- Charness, Gary, "Self-Serving Cheap Talk: A Test of Aumann's Conjecture," *Games and Economic Behavior*, XXXIII (2000), 177–194.
- Chwe, Michael Suk-Young, "Farsighted Coalitional Stability," *Journal of Economic Theory*, LXIII (1994), 299–325.
- Cooper, Russell, Douglas V. De Jong, Robert Forsythe, and Thomas W. Ross, "Communication in Coordination Games," *Quarterly Journal of Economics*, CVII (1992), 739–771.
- Dagan, Nir, Roberto Serrano, and Oscar Volij, "A Noncooperative View of Consistent Bankruptcy Rules," *Games and Economic Behavior*, XVIII (1997), 55–72.
- Delgado, Juan, and Diego Moreno, "Coalition-Proof Supply Function Equilibria in Oligopoly," *Journal of Economic Theory*, CXIV (2004), 231–254.
- Farrell, Joseph, "Communication, Coordination and Nash Equilibrium," *Economics Letters*, XXVII (1988), 209–214.
- Fudenberg, Drew, and Jean Tirole, *Game Theory*, 2nd edition (Cambridge, MA: MIT Press, 1992).
- Gul, Faruk, "Rationality and Coherent Theories of Rational Behavior," *Journal of Economic Theory*, LXX (1996), 1–31.
- Gupta, Srabana, and Richard E. Romano, "Monitoring the Principal with Multiple Agents," *Rand Journal of Economics*, XXIX (1998), 427–442.
- Konishi, Hideo, Michel LeBreton, and Shlomo Weber, "On Coalition-Proof Nash Equilibria in Common Agency Games," *Journal of Economic Theory*, LXXXV (1999), 122–139.
- Mariotti, Marco, "A Model of Agreements in Strategic Form Games," *Journal of Economic Theory*, LXXIV (1997), 196–217.
- Noe, Thomas H., "Rationalizable and Coalition Proof Shareholder Tendering Strategies in Corporate Takeovers," *Review of Quantitative Finance & Accounting*, XI (1998), 269–291.
- Pearce, David G., "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica*, LII (1984), 1029–1050.
- Rabin, Matthew, "Communication between Rational Agents," *Journal of Economic Theory*, LI (1990), 144–170.
- , "Incorporating Behavioral Assumptions into Game Theory," in James Friedman, ed., *Problems of Coordination in Economic Activity* (Boston and Dordrecht: Kluwer Academic, 1994b), pp. 69–87.
- Ray, Debraj, and Rajiv Vohra, "Equilibrium Binding Agreements," *Journal of Economic Theory*, LXXIII (1997), 30–78.
- Ray, Debraj, and Rajiv Vohra, "A Theory of Endogenous Coalition Structures," *Games and Economic Behavior*, XXVI (1999), 286–336.
- Schelling, Thomas C., *The Strategy of Conflict* (Cambridge, MA: Harvard University Press, 1960).
- Tan, Tommy Chin-Chiu, and Sergio Ribeiro da Costa Werlang, "The Bayesian Foundations of Solution Concepts of Games," *Journal of Economic Theory*, XVLV (1988), 370–391.
- Watson, Joel, "Communication and Superior Cooperation in Two-Player Normal Form Games," *Economics Letters*, XXXV (1991), 267–271.
- Xue, Licun, "Coalitional Stability under Perfect Foresight," *Economic Theory*, XI (1998), 603–627.
- , "Negotiation-Proof Nash Equilibrium," *International Journal of Game Theory*, XXIX (2000), 339–357.