Theoretical Economics 15 (2020), 669–714 1555-7561/20200669

Equilibrium coalitional behavior

Mert Kimya
School of Economics, University of Sydney

I develop two related solution concepts—equilibrium coalitional behavior and credible equilibrium coalitional behavior—that capture foresight and impose the requirement that each coalition in a sequence of coalitional moves chooses optimally among all its available options. The model does not require, but may use, the apparatus of a dynamic process or a protocol that specifies the negotiation procedure underlying coalition formation. Therefore, it forms a bridge between the non-cooperative and the cooperative approaches to foresight.

Keywords. Coalition formation, farsightedness.

JEL classification. C70, C71, C72, D71.

1. Introduction

This paper contributes to the literature on farsighted coalition formation. I define a cooperative domain—extended coalitional games—and two related solution concepts on this domain—equilibrium coalitional behavior (ECB) and credible equilibrium coalitional behavior (CECB)—that capture foresight.

Extended coalitional games

Let N be a set of players and let Z be a set of nodes. Let A_z denote the set of all possible coalitional actions available at node z, where an action is a triple (z, z', S) that denotes the possibility of coalition S to move the game from node z to node z'. At some nodes, it might not be feasible for any coalition to take an action. This is represented as the particular action $(z, z, \varnothing)$ that denotes remaining at node z. A path is a sequence of *feasible* actions. A terminal path is a path that is either infinite or that ends with no action.

An *extended coalitional game* is defined as $\Gamma = \{N, Z, \{A_z\}_{z\in Z}, \{\succeq_i\}_{i\in N}\}$, where the preferences are defined over the set of terminal paths.

An extended coalitional game is a direct generalization of an extensive form game of perfect information, but it is most closely related to the domain of the abstract game

Mert Kimya: mert.kimya@sydney.edu.au
I am particularly grateful to Rajiv Vohra for continuous support and guidance throughout the preparation of this work. I would also like to thank three anonymous referees, Pedro Dal Bo, Mark Dean, Geoffroy de Clippel, Elif Cansu Dilmen, Jack Fanning, Roberto Serrano, Licun Xue, seminar participants at Bogazici University, Istanbul Technical University, University of Essex, University of Sussex, University of Sydney, X. Winter Workshop at Koç University, 22nd Coalition Theory Workshop, and SAET 2018 for helpful comments and suggestions.

Available at https://econtheory.org. https://doi.org/10.3982/TE3055

(see, for example, Chwe 1994, Rosenthal 1972, and Xue 1998). There are two differences between an abstract game and an extended coalitional game.

(i) In an abstract game, utilities are defined over the nodes, whereas in an extended coalitional game, they are defined over the paths. This is the major difference between the two domains. Nevertheless, given an abstract game, the utility a player gets from a path can be appropriately defined based on the approach one wants to use. There are two main approaches in the literature: the static approach and the dynamic approach. The former assumes that players care only about the final outcome the negotiations lead to; hence, the utility of a path should correspond to the utility of the final node on the path. The latter approach assumes that players discount the utilities with a discount factor δ; hence, the utility of a path should correspond to the discounted utility of the nodes along the path (see Section 4.2 for a formal analysis).[1]

The point is that defining the utilities over the paths cause no loss in generality. This is because the approaches on the abstract game also implicitly use utilities over paths derived from the utilities on the nodes.

(ii) In an abstract game, it is always possible to take no action at each node, whereas in an extended coalitional game the modeler is free to choose whether to include no-action as a possible action. This is a minor difference that is self-evident.

(Credible) equilibrium coalitional behavior

A *coalitional behavior* is a complete plan of action defined on the extended coalitional game. That is, it assigns a unique action to each node of an extended coalitional game. Thereby, it also assigns a "path of play" (terminal path) to each node of an extended coalitional game.

A coalition $S \subseteq N$ can deviate from a coalitional behavior by refusing to take some of the prescribed actions, for this S needs to have a nonempty intersection with the coalitions taking these actions. Instead, S can take actions unspecified by the coalitional behavior, for this S needs to contain the coalitions that have the power to take these actions. A deviation is profitable if every $i \in S$ prefers the resulting path of play to the initial path of play at each node where an action changes.

An ECB is simply a coalitional behavior that is immune to profitable deviations.

A profitable deviation is credible if it cannot be followed by further deviations that would make the initial deviator worse off, in a sense that is clarified later. A CECB is a coalitional behavior that is immune to profitable and credible deviations.

1.1 *Contribution*

The literature on farsighted coalition formation has developed extensively in recent years. Earlier solution concepts suffered from what is known as the problem of maximality (see Ray and Vohra 2014), where players form unreasonable expectations instead of taking the best course of action available to them.

[1]There are other approaches. For instance, some authors define a non-cooperative game from the abstract game and define their concepts over this non-cooperative game. Notable examples include Herings et al. (2004) and Granot and Hanany (2016). This approach is discussed in Section 4.2.1.

(C)ECB directly overcomes this problem by imposing consistent expectations. This is not new as several other solution concepts have overcome this problem similarly before (see Dutta and Vohra 2017, Konishi and Ray 2003, and Ray and Vohra 2014). There are two ways in which (C)ECB supports this claim further.

Backward Induction. Intuitively, one might expect a farsighted solution concept to incorporate considerations similar to backward induction in a cooperative setting. Nevertheless, in the absence of the structure of an extensive form, it is not clear how to incorporate such considerations. Extended coalitional games have this structure in certain contexts and indeed we see that (C)ECB can be found through backward induction whenever it is possible.

Non-cooperative Justification. The problem of maximality is not present in non-cooperative solution concepts such as the subgame perfect equilibrium. Hence, if it is possible to develop an intuitive non-cooperative game from an extended coalitional game, whose subgame perfect equilibrium outcomes coincide with the cooperative solution concept, then this is a good indication that the solution concept we start with incorporates the notion of maximality. With (C)ECB, this exercise can be done under certain conditions on the domain.

Second, (C)ECB directly looks at possible profitable deviations to see if a plan of action is stable. Some other solution concepts, such as the equilibrium process of coalition formation (see Konishi and Ray 2003 and Ray and Vohra 2014) and the (strong) rational expectations farsighted stable set (see Dutta and Vohra 2017), also require immunity from deviations. A *process of coalition formation* is a dynamic process that specifies the coalition that moves at each node and the particular node the coalition moves to. An *equilibrium process of coalition formation* (EPCF) is a process of coalition formation, where coalitional moves at every state are immune to single-step deviations and preferable to inaction. *(Strong) rational expectations farsighted stable set* ((S)REFS) is a set of nodes that satisfy certain farsighted external and internal stability conditions and that can be supported by an expectation function, which specifies the transitions between nodes. Although both EPCF and (S)REFS also require immunity from deviations, (C)ECB diverges in one important aspect: it does not restrict attention to one-step deviations.

One might surmise that restriction to one-step deviations is without loss of generality, but depending on the situation and how the utilities are defined, this might not be the case. Indeed, there are situations where one cannot replicate a profitable deviation by a one-step deviation under (C)ECB. Mainly because, a deviation to a "cycle" (or an infinite path) might be the only profitable deviation available.

Hence, the issue is also related to the issue of cycles. The static solution concepts never prescribe cycles. As the utilities are defined over the paths, (C)ECB allows cycles to be incorporated as in the dynamic approach. But even the dynamic solution concepts such as the EPCF ignore cycles to some extent by restricting attention to one-step deviations. In some situations this assumption rules out (possibly profitable) deviations to cycles (see the example in Figure 7 in Section 4.2). This is no longer the case with (C)ECB.

Finally, and perhaps most importantly, the literature on farsighted coalition formation is very fragmented and (C)ECB shows that it is possible to connect the different approaches to foresight.

(C)ECB is a versatile solution concept that might be used to study a variety of domains existing in the literature. The versatility of (C)ECB is a result of the generality of its domain: extended coalitional games.

In an extended coalitional game, utilities of the players are defined over the paths. This allows one the flexibility to define the solution concept as a static concept, where players care only about the final outcome the negotiations lead to, and as a dynamic concept, where players also care about the outcomes along the paths through which an agreement is reached.

In this way (C)ECB is able to bridge these two different strands of the literature on farsighted coalition formation. Indeed, ECB can be directly related to solution concepts in both approaches: SREFS in the static approach and EPCF in the dynamic approach.[2]

The versatility of the extended coalitional games also allow us to study certain non-cooperative games with (C)ECB. In particular, extensive form games of perfect information are extended coalitional games and (C)ECB is closely related to the most popular solution concept on this domain: the subgame perfect equilibrium. In finite extensive form games of perfect information, (C)ECB reduces down to subgame perfect equilibrium and ECB refines subgame perfect equilibrium in infinite horizon games.

The variety of the approaches and solution concepts used to study farsighted coalition formation is one of the problems of the literature: "as one surveys the landscape of this area of research, the first feature that attracts attention is the fragmented nature of the literature" Ray and Vohra (2014).

Hence, it would be useful to have a framework that could connect the different approaches to foresight. By allowing a translation of the assumptions of the different approaches into its own structure, extended coalitional games provide this framework to a certain extent.

Indeed, ECB shows that the developments in the literature have made it possible to achieve some level of unification between both the non-cooperative and cooperative strands of literature and within the cooperative approach. The evidence for this is the relationships ECB engenders with solution concepts in all three approaches discussed above. This is especially important when one considers how fragmented the literature on farsighted coalition formation is.

1.2 *Outline of the results*

A CECB exists in finite extended coalitional games. In finite acyclic games,[3] both ECB and CECB satisfy the *one-step deviation property* and, hence, can be found through a simple backward induction algorithm. Also, an algorithm to find the CECB can be provided for finite games with cycles.

[2] CECB is also related to SREFS (albeit not as strongly as ECB), but not to EPCF.

[3] A game is acyclic if whenever node $z \neq z'$ is reachable from z' through some sequence of actions, z' is not reachable from z.

On the non-cooperative side, one can find a simple bargaining game defined over an extended coalitional game whose stationary subgame perfect equilibrium is related to (C)ECB. In particular, in games that satisfy a condition called strong acyclicity, each ECB can be supported as a stationary subgame perfect equilibrium of the bargaining game. Furthermore, if we impose further conditions, then we can obtain full implementation of both ECB and CECB through the stationary subgame perfect equilibrium of the bargaining game.

One attractive feature of (C)ECB is that it can be applied to general classes of games that are used extensively in the literature. In Section 4, I choose three such domains: the extensive form games, the abstract game, and characteristic function games. I demonstrate that (C)ECB is related to attractive solution concepts in these domains: subgame perfect equilibrium in extensive form games, SREFS and EPCF in the abstract game, and the core in characteristic function games.

Finally, to demonstrate the strength and novelty of the solution concepts, I apply the (C)ECB to network formation games and show that it is able to make novel predictions in this widely studied environment. One of the findings of the strategic network formation literature is the tension between efficiency and stability. Furthermore, Dutta et al. (2005) and Herings et al. (2009) show that this tension persists even if we consider farsighted solution concepts. It turns out that (C)ECB is able to overturn these results. In particular, we can show that there is a "reasonable" way to allocate the value of a network such that every efficient graph can be supported as the prediction of a (C)ECB when players are sufficiently patient. I also provide conditions under which all of the predictions of a (C)ECB would be efficient and uncover a relationship between the solution concept and what I call the pessimistic network core.

I start with the description of the domain and the solution concept in Section 2. Section 3 analyzes the properties of (C)ECB, Section 4 studies (C)ECB and compares it to solution concepts in extensive form games, characteristic function games, and the abstract game. Section 5 studies network formation under (C)ECB, Section 6 includes the literature review, and Section 7 concludes with a remark. All proofs are provided in the Appendix.

2. Preliminaries

2.1 *Extended coalitional games*

An *extended coalitional game* is defined as $\Gamma = \{N, Z, \{A_z\}_{z\in Z}, \{\succeq_i\}_{i\in N}\}$, where N is the finite set of players, Z is the set of nodes, A_z is the set of actions available at node $z \in Z$, and $\succeq_i$ is the preference relation of player $i \in N$ over the set of terminal paths, which are defined shortly.

An action is a triple (z, z', S), where the first entry $z \in Z$ denotes the node at which the action can be taken, the second entry $z' \in Z$ denotes the node to which the action is leading, and the third entry denotes the coalition $S \subseteq N$ that can take the action. This coalition is called the *initiator* of the action. At some nodes, it might be possible for no coalition to take an action. This is represented as the particular action $(z, z, \varnothing)$ that denotes the possibility of remaining at node z. The set of actions A_z is required to be nonempty: $A_z = \{(z, z, \varnothing)\}$ at a no-action node z.

A path is a sequence of actions $\{a_k\}_{k=1,\dots,K} = \{(z_k, z_{k+1}, S_k)\}_{k=1,\dots,K}$, where K might be infinite. A path $\{a_k\}_{k=1,\dots,K}$ is terminal if it is infinite or if $a_K = (z_K, z_K, \varnothing)$. Let $\mathcal{P}$ denote the set of all terminal paths and let $\succeq_i$ denotes the preference relation of $i \in N$ on $\mathcal{P}$.

2.2 *Equilibrium coalitional behavior*

A coalitional behavior is a complete plan of action. It prescribes a unique action to each node of an extended coalitional game. Let A denote the set of all actions, i.e., $A = \bigcup_{z\in Z} A_z$.

Definition 1 (Coalitional behavior). A coalitional behavior is a mapping $\phi : Z \to A$, where $\phi(z) \in A_z$ for all $z \in Z$.

In turn, each coalitional behavior also prescribes a unique terminal path to each node of an extended coalitional game; this path is called the path of play. Let $\mathcal{P}_z$ denote the set of terminal paths that start at node z.

Definition 2 (Path of play). Given a coalitional behavior ϕ, a path of play is a mapping $\sigma : Z \to \mathcal{P}$ such that for each $z \in Z$, $\sigma(z) = \{(z_k, z_{k+1}, S_k)\}_{k=1,\dots,K} = \{\phi(z_k)\}_{k=1,\dots,K} \in \mathcal{P}_z$, where $z_1 = z$.

Throughout the paper, when the coalitional behavior is denoted by ϕ, the corresponding path of play is denoted by σ, and when the coalitional behavior is denoted by ϕ', the corresponding path of play is denoted by σ'. To define our solution concept, we need the definition of a coalitional deviation from a prescribed plan of action. Intuitively, no coalition can be forced to take an action. Hence, each coalition S might refuse to take a prescribed action $\phi(z)$ if S has a nonempty intersection with the initiators of $\phi(z)$. Instead, S can take an action (z, z', T), which it has the power to implement, i.e., where S contains the initiators of this action. A deviation by S is profitable if at every node at which an action changes, everybody in S prefers the new path of play to the initially prescribed path of play.

Definition 3 (Coalitional deviation). The relationship $S \subseteq N$ has a deviation from a coalitional behavior ϕ to a coalitional behavior ϕ' if for every $z \in Z$ such that $\phi(z) \neq \phi'(z)$ we have the following relationships.

- If $\phi(z) = (z, z', T)$, where $T \neq \varnothing$, then $S \cap T \neq \varnothing$. (If an action specified by ϕ is not taken, then S has a member who can refuse to take this action.)
- If $\phi'(z) = (z, z', T)$, then $S \supseteq T$. (If an action not specified by ϕ is taken, then S should be able to induce this action.)

We say that the deviation by S is profitable if for every $z \in Z$ such that $\phi(z) \neq \phi'(z)$, we have $\sigma'(z) \succ_i \sigma(z)$ for all $i \in S$.

Note that the definition allows for a coalition $T \subseteq S$ to block a specified action (z, z', S) and take no other action at z, if no action is possible at z. An ECB is simply a coalitional behavior that is immune to profitable deviations.

Definition 4 (Equilibrium coalitional behavior). A coalitional behavior ϕ is an ECB if there does not exist a profitable coalitional deviation from ϕ.

Note the simplicity of the definition, which is in contrast to most other farsighted solution concepts. The typical approach in the literature is to look at a set of outcomes or paths that has some "stability" properties, whereas ECB directly looks at profitable deviations from plans of actions.

In contrast to this simplicity, ECB is also a powerful concept in the sense that it is related to a wide variety of solution concepts from different strands of the literature, such as subgame perfect equilibrium, the core, SREFS and EPCF. However, ECB has issues regarding existence. CECB gets over this problem by also considering the credibility of a profitable deviation.

2.3 *Credible equilibrium coalitional behavior*

The subgame at $z \in Z$, denoted by $\Gamma(z)$, is the game that includes only those nodes that are reachable from z and the actions between these nodes. The utilities are defined the same way as in the original game. We say that a subgame is nontrivial if it includes (strictly) more than one node.

Definition 5 (Basic game). An extended coalitional game is a basic game if the only nontrivial subgame of the game is the game itself.[4]

An example of a basic game is a tree of length 1; another example would be a standalone cycle (see Figure 1(b) and Figure 2(b), (c) for examples of basic games). We first define the notion of credibility for a basic game and then extend it to any extended coalitional game.

The idea is that a profitable deviation should only be credible if it cannot be followed by further farsighted deviations that would make one of the initial deviators worse off. To incorporate this idea, I define a credible set of coalitional behaviors as a set of coalitional behaviors such that any profitable deviation from a coalitional behavior in the set might be followed by further farsighted profitable deviations back into a coalitional behavior in the credible set that makes one of the initial deviators worse off. Hence, any profitable deviation from a credible coalitional behavior is avoided by further deviations to a credible coalitional behavior.

Let Φ denote the set of all coalitional behaviors. For $\phi, \phi' \in \Phi$, we say that ϕ dominates ϕ', $\phi >_D \phi'$, if there exists a profitable deviation from ϕ' to ϕ. We write $\phi >_D^S \phi'$ if there exists a profitable deviation from ϕ' to ϕ by coalition $S \subseteq N$. Finally, for any

[4]Formally, Γ is a basic game if for every $z \in \Gamma$, either $\Gamma(z) = \Gamma$ or $z' \in \Gamma(z)$ implies that $z' = z$.

two coalitional behaviors $\phi, \phi' \in \Phi$, $\phi \succ_i \phi'$ if for every $z \in Z$ such that $\phi(z) \neq \phi'(z)$, $\sigma(z) \succ_i \sigma'(z)$.

The definition of indirect dominance through profitable deviations below is the usual indirect dominance relation due to Harsanyi (1974) and Chwe (1994), where the relation is defined over the set of coalitional behaviors.[5]

Definition 6 (Indirect dominance through profitable deviations). A coalitional behavior $\phi \in \Phi$ indirectly dominates $\phi' \in \Phi$ through profitable deviations, denoted by $\phi \rhd \phi'$, if there exists $\phi_0, \phi_1, \dots, \phi_n \in \Phi$ and $S_0, S_1, S_2, \dots, S_{n-1}$ such that $\phi_0 = \phi'$, $\phi_n = \phi$, $\phi_{i+1} >_D^{S_i} \phi_i$, and $\phi \succ_{S_i} \phi_i$ for all $i = 0, \dots, n-1$.

A set of coalitional behaviors Y is a credible set if a coalitional behavior ϕ is in Y if and only if any profitable deviation from ϕ can be followed by a sequence of profitable deviations into Y, which makes one of the initial deviators worse off. Furthermore, the sequence of profitable deviations has the property that each coalition moving along the sequence foresees the final coalitional behavior the sequence leads to and prefers that final coalitional behavior to the coalitional behavior it replaces.

Definition 7. In a basic game, a set of coalitional behaviors $Y \subseteq \Phi$ is credible if $\phi \in Y$ if and only if for all $S \subseteq N$ and $\phi' \in \Phi$ such that $\phi' >_D^S \phi$ there exists $\phi^* \in Y$ with $\phi^* \rhd \phi'$ and $\phi^* \not\succ_S \phi$.[6]

The definition of a credible set mirrors the definition of a consistent set of Chwe (1994) with two major differences: (a) the focus is on coalitional behaviors instead of the outcomes and (b) the profitable deviation relation is taken as the effectiveness relation in Chwe (1994).[7]

There might be multiple credible sets, but the lemma below, which is due to Chwe (1994), establishes that there always exists a unique credible set that contains all of the other credible sets. This set is called the *largest credible set* (LCRS), following the largest consistent set of Chwe (1994).

Lemma 1. *In any basic game, there uniquely exists a Y such that Y is credible and (Y' credible $\implies$ $Y' \subseteq Y$). The Y is called the LCRS.*

In a basic game, we say that a coalitional behavior is credible if it is included in the LCRS. Similarly, we say that a profitable deviation from ϕ to ϕ' by coalition S is credible if it cannot be followed by profitable deviations to a credible coalitional behavior ϕ^* such that $\phi^* \not\succ_S \phi$.

[5]See Section 4.2 and Definition 14 for the standard definition of an indirect dominance relation in an abstract game.

[6]We have $\phi^* \succ_S \phi$ if $\phi^* \succ_i \phi$ for all $i \in S$.

[7]Point (b) is reminiscent of the difference between the original definition of indirect dominance given in Harsanyi (1974) and the modified definition in Chwe (1994). The former imposes that the coalition improves at every step. In this sense, the definition used here is closer to Harsanyi's (1974) definition. As you will see, the current definition has the advantage of making ECB a refinement of CECB.

Definition 8 (Credible deviation in a basic game). In a basic game, a profitable deviation from ϕ to ϕ' by coalition S is credible if there does not exist a credible coalitional behavior ϕ^* such that $\phi^* \rhd \phi'$ and $\phi^* \not\succ_S \phi$.

It is easy to see that, in a basic game, a coalitional behavior ϕ is immune to profitable and credible deviations if and only if ϕ is in the LCRS. One could choose to use the LCRS as the solution concept in the whole extended coalitional game instead of basic games. But LCRS has some problems that I would like to avoid; see Remark 1 for a brief discussion of these problems and see Section 7 for a more detailed explanation. To avoid these problems, I extend the definition of a credible deviation from a basic game to any extended coalitional game using the notion of "reduced games," which is explained now.

Once a coalition S deviates from a coalitional behavior ϕ' to a coalitional behavior ϕ, the players would expect the game to continue as in ϕ at any proper subgame $\Gamma(z)$. That is, players might assume that any action leading to the proper subgame $\Gamma(z)$ is going to be followed by $\sigma(z)$. I define a reduced game based on this idea.

We say that z *leads* to a proper subgame of $\Gamma(z^*)$ if $\Gamma(z)$ is a proper subgame of $\Gamma(z^*)$ and if z is directly reachable from a node that shares the same subgame as z^*. Let $\underline{z^*}$ denote the set of nodes that lead to a proper subgame of $\Gamma(z^*)$. Formally, $\underline{z^*} = \{z \in \Gamma(z^*) \mid \Gamma(z) \subset \Gamma(z^*)$ and there exists $(z', z, S) \in A$ such that $\Gamma(z^*) = \Gamma(z')\}$.

If $\Gamma(z^*)$ is a tree, then $\underline{z^*}$ is composed of those nodes that are directly reachable from z^*. For instance, in the game depicted in Figure 1, $\underline{a} = \{b, c, d\}$. If z^* is on a cycle, then $\underline{z^*}$ is composed of those nodes that are directly reachable from the cycle that z^* belongs to, but it does not include any node that is in the same cycle with z^*. For instance, in the game depicted in Figure 2, $\underline{a} = \underline{b} = \{c\}$.

To see whether a deviation is credible at node z^*, the players can simply "remove" the subgames that start at nodes in $\underline{z^*}$ and treat those nodes in $\underline{z^*}$ as terminal nodes, where the payoff of any node in $\underline{z^*}$ is given by the associated path of play. The definition of a reduced game formalizes this idea.

Definition 9 (Reduced game). Given the game $\Gamma = \{N, Z, \{A_z\}_{z\in Z}, \{\succeq_i\}_{i\in N}\}$, the reduced game at $z^* \in Z$ given ϕ, denoted by $\Gamma(z^*, \phi)$, is defined as $\Gamma(z^*, \phi) = \{N, Z^*, \{A^*_z\}_{z\in Z^*}, \{\succeq^*_i\}_{i\in N}\}$, where the following statements hold:

- We have $Z^* = \{z \in Z \mid$ either $z \in \underline{z^*}$ or $(z \in \Gamma(z^*)$ and $z \notin \Gamma(z')$ for any $z' \in \underline{z^*})\}$. (All nodes that are part of a proper subgame are removed from $\Gamma(z^*)$ except those leading to a proper subgame.)
- We have $A^*_z = A_z$ if $z \in Z^* \setminus \underline{z^*}$ and $A^*_z = (z, z, \varnothing)$ if $z \in \underline{z^*}$. (Nodes that lead to a proper subgame are treated as terminal nodes.)
- For any terminal path p, let $p^* = (p, \sigma(z))$ if p terminates at $z \in \underline{z^*}$; otherwise let $p^* = p$. Then for any two terminal paths p_1, p_2 and for any $i \in N$, we have $p_1 \succeq^*_i p_2$ if and only if $p^*_1 \succeq_i p^*_2$. (Each path p terminating at a node z leading to a proper subgame has the "utility" of the terminal path $(p, \sigma(z))$ in the original game.)

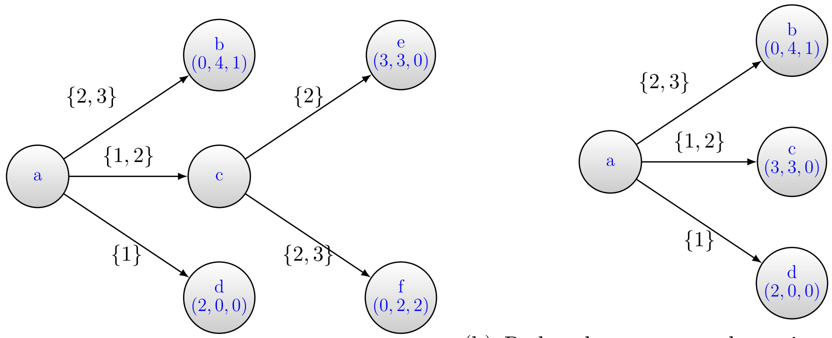

(a) The game

(b) Reduced game at node a given ϕ_2, where $\phi_2(c) = (c, e, \{2\})$

Figure 1. An acyclic game.

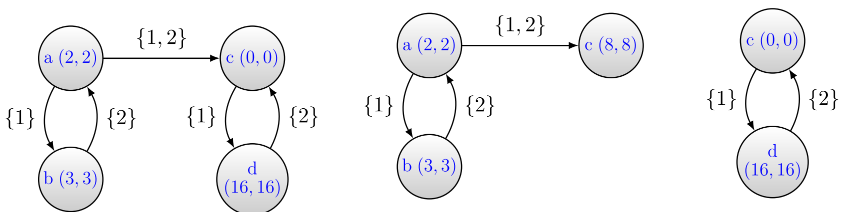

(a) The game

(b) Reduced game at node a given ϕ_1

(c) Reduced game at node c

Figure 2. A game with cycles.

The main idea behind a reduced game is simple: at node z^*, the players expect the game to be continued as in ϕ in any proper subgame of $\Gamma(z^*)$. Hence, they treat each node z leading to a proper subgame as a terminal node and they assign the utility of the terminal path $(p, \sigma(z))$ to any path p terminating at z. One important point to note is that while finding reduced games, we never break cycles. This is because reduced games are obtained by removing proper subgames and two nodes in a cycle are always associated with the same subgame. Finally, note that each reduced game is necessarily a basic game.

See Figure 1 for a reduced game in a tree structure. Assume that the utility of a terminal path is the utility of the node the path terminates in and consider the coalitional behavior that assigns $\phi(c) = (c, e, \{2\})$. To find the reduced game at the root, note that there is only one (nontrivial) proper subgame, which starts at c. Once we remove every node reachable from c and assign c the utility associated with the path prescribed by the coalitional behavior ϕ, we are left with the game in Figure 1(b).

See Figure 2 for an example with cycles. Assume that no-action is available at every node and the utility of a terminal path is the discounted utility of the nodes along the path with discount factor $\delta = \frac{1}{2}$. Consider the coalitional behavior ϕ that assigns $\phi(c) = (c, d, \{1\})$ and $\phi(d) = (d, d, \varnothing)$. To find the reduced game at node a, observe that the game has only one proper subgame, which is the cycle composed of c and d. Hence, we remove only this cycle. We treat node c, which connects this cycle to the game, as a terminal node and assign the utility associated with $\sigma(c)$ to this node. In this case, ϕ specifies that at node c, the players will receive a payoff of 16 indefinitely but with one period delay. Since $\delta = \frac{1}{2}$, this is equivalent to receiving a payoff of 8 indefinitely, so we can plug in $(8, 8)$ on node c. Finally, we obtain the reduced game in Figure 2(b).

We have seen the reduced game at node a, but what about the reduced game at any node on the cycle composed of c and d? The subgame at node c and node d is the cycle itself and this cycle does not contain any proper subgames. Hence, the reduced game at any node in the cycle would be the cycle itself. For instance, $\Gamma(c, \phi) = \Gamma(c)$ for any coalitional behavior ϕ (see Figure 2(c)).

Each reduced game is necessarily a basic game. Hence, we can check the credibility of a deviation at each reduced game, the idea being that if the players expect the game to be continued as in ϕ, then at any reduced game their deviation should be credible.

To formalize this idea, we need an additional definition. The deviation from ϕ' to ϕ restricted to the reduced game $\Gamma(z, \phi)$ is the deviation in $\Gamma(z, \phi)$ from the coalitional behavior that matches with ϕ' in the nonterminal nodes included in $\Gamma(z, \phi)$ to the coalitional behavior that matches with ϕ in the nonterminal nodes included in $\Gamma(z, \phi)$.

Definition 10 (Credible deviation). A profitable deviation from ϕ' to ϕ is credible if for each $z \in Z$ for which $\phi(z) \neq \phi'(z)$, the deviation restricted to $\Gamma(z, \phi)$ is credible in the reduced game $\Gamma(z, \phi)$.

Definition 11 (Credible ECB). A coalitional behavior ϕ is a CECB if it is immune to profitable and credible deviations.

In a basic game, ϕ is a CECB if and only if ϕ is an element of the LCRS. The use of reduced games to define CECB implies that this equivalence no longer holds when we go beyond basic games. Remark 1 below explains why we use the formalism of reduced games instead of directly looking at the LCRS of the whole game.

Remark 1 (LCRS versus CECB). Imagine a situation, where the game comes to a point in which a player chooses between a payoff of 1 and a payoff of 2, and when the player makes this decision the game ends. In such a game, LCRS might as well specify that the player would choose the lower payoff because she can be unreasonably concerned that her deviation to the higher payoff might trigger a wave of deviations in the past that could make her worse off. Such an example is formalized in Section 7.

With the use of reduced games, CECB overcomes such problems. This provides CECB with certain desirable properties such as the fact that it can be found with backward induction whenever possible. Furthermore LCRS is a very permissive solution concept, and the reduced game formalism allows CECB to alleviate this multiplicity problem to a great extent without compromising it from existence. All of these points are shown with a concrete example in Section 7.

2.4 *Examples*

Consider the game represented in Figure 1(a). There are three players, 1, 2, and 3. There are six nodes, four of which are terminal. The utility of a terminal path is the utility of the node the path terminates in.

Consider the coalitional behavior ϕ_1 that assigns $\phi_1(a) = (a, d, \{1\})$ and $\phi_1(c) = (c, f, \{2, 3\})$. This is not an ECB, because the coalition $\{1, 2\}$ has a profitable deviation in which player 1 blocks the action $(a, d, \{1\})$. Instead the coalition $\{1, 2\}$ takes the action $(a, c, \{1, 2\})$ and player 2 blocks the action $(c, f, \{2, 3\})$ at node c, and, instead, takes the action $(c, e, \{2\})$. Let us call the resulting coalitional behavior ϕ_2, where $\phi_2(a) = (a, c, \{1, 2\})$ and $\phi_2(c) = (c, e, \{2\})$.

But now the coalition $\{2, 3\}$ has a profitable deviation from ϕ_2 in which player 2 blocks the action $(a, c, \{1, 2\})$ and, instead, the coalition $\{2, 3\}$ takes the action $(a, b, \{2, 3\})$. Let us call the resulting coalitional behavior ϕ_3, where $\phi_3(a) = (a, b, \{2, 3\})$ and $\phi_3(c) = (c, e, \{2\})$. It is easy to see that there is no profitable deviation from ϕ_3 and, indeed, this is the unique ECB of this game.

CECB incorporates the idea that the initial deviation from ϕ_1 is not credible since it can be followed by another deviation that would make the initial deviator worse off. To see this formally, consider the reduced game at node a given ϕ_2, which is depicted in Figure 1(b). The unique LCRS of this game is $\{(a, b, \{2, 3\}), (a, d, \{1\})\}$. Note that $(a, d, \{1\})$ is in the LCRS, because any profitable deviation will be followed by another profitable deviation to an action in the LCRS that makes one of the initial deviators (player 1) worse off. Hence, the initial profitable deviation is not credible.

Nevertheless, there exists a profitable and credible deviation from ϕ_1 in which player 2 blocks the action $(c, f, \{2, 3\})$ and, instead, takes the action $(c, e, \{2\})$. Note that the reduced game at node c is the subgame that starts at node c, which has a singleton LCRS composed of $(c, e, \{2\})$; hence, the deviation is credible. Let us call the resulting coalitional behavior ϕ_4, where $\phi_4(a) = (a, d, \{1\})$ and $\phi_4(c) = (c, e, \{2\})$.

The behavior ϕ_4 is a CECB and since ϕ_3 is an ECB, it is also a CECB. One can easily check that these are all the CECBs of this game.

For an example with cycles, look at the game in Figure 2(a). Assume that no-action is available at every node and the utility of a terminal path is the discounted utility of the nodes along the path with discount factor $\delta = \frac{1}{2}$, i.e., the utility of player $i \in N$ from the terminal path $\{(z_k, z_{k+1}, S_k)\}_{k=1,2,\dots,K}$ is $\sum_{k=1,\dots,K} \delta^{k-1} u_i(z_k) + \sum_{k=K,\dots} \delta^k u_i(z_K)$, where $\delta = \frac{1}{2}$.

Consider the subgame at node c (see Figure 2(c)), which is a basic game. In this game, LCRS is composed of a single coalitional behavior ϕ_1, which assigns $\phi_1(c) = (c, d, \{1\})$ and $\phi_1(d) = (d, d, \varnothing)$. To see this, observe that in this subgame there is a profitable deviation from any other coalitional behavior to ϕ_1.

But then, in the whole game, any CECB ϕ should necessarily assign $\phi(c) = (c, d, \{1\})$ and $\phi(d) = (d, d, \varnothing)$. Otherwise there is a deviation from the coalitional behavior, in which only the actions at node c and d change to the actions assigned by ϕ_1. Since only these actions change, and since the LCRS of the reduced game at both node c and node d is composed of ϕ_1, we have that the deviation is credible and profitable.

Now observe that the LCRS of the reduced game at node a given ϕ_1 (see Figure 2(b)) is composed of a single coalitional behavior ϕ_2, which assigns $\phi_2(a) = (a, c, \{1, 2\})$ and $\phi_2(b) = (b, a, \{2\})$. This is because there is a profitable deviation from any other coalitional behavior to ϕ_2 and ϕ_2 is immune to profitable deviations.

When we combine ϕ_1 and ϕ_2, we obtain a coalitional behavior that is necessarily a CECB for the whole game. Let us call it ϕ. If there exists any other CECB ϕ', by the observation above, $\phi'(c) = \phi(c)$ and $\phi'(d) = \phi(d)$. If $\phi'(j) \neq \phi(j)$ for $j = a, b$, then there will be a profitable deviation to ϕ. Furthermore, the deviation will necessarily be credible because the reduced game given ϕ (depicted in Figure 2(b)) has a singleton LCRS composed of ϕ_2, which coincides with ϕ in this reduced game.

This implies that the unique CECB of the game is ϕ. Since ϕ is immune to profitable deviations, it is also the unique ECB.[8]

An easier way to find the (C)ECBs in both of these examples is to use backward induction. That is, one can solve these games by recursively finding (C)ECBs of basic games, starting from subgames that are themselves basic games and then plugging in the utilities of the found path of plays onto the nodes *leading* to these subgames. One can continue doing this until every (C)ECB of the game is found. This is shown rigorously in Section 3.3, which also provides an algorithm to find the CECBs in finite cyclic games.

3. Properties

3.1 *Existence*

Proposition 1. *CECBs exists in any finite extended coalitional game.*

The existence of CECBs in basic games reduces down to the existence and nonemptiness of the LCRS in basic games. Chwe (1994) shows that a largest consistent set would exist and be nonempty. This directly implies the existence and nonemptiness of the LCRS in basic games, hence the existence of CECB in basic games.

In any finite game, one can find the CECBs through recursively finding the CECBs of basic games. Hence, a CECB exists in any finite game. See the Appendix for details.

3.2 *One-step deviation property*

We say that a deviation is a one-step deviation if every action involved in the deviation stems from the same node. We say that ECB satisfies the one-step deviation property if the existence of a profitable deviation implies the existence of a profitable one-step deviation. Similarly, CECB satisfies the one-step deviation property if the existence of a credible and profitable deviation implies the existence of a one-step credible and profitable deviation. An acyclic game is a game that does not contain cycles.

Proposition 2. *In finite acyclic games, both ECB and CECB satisfy the one-step deviation property.*

[8]This is because each ECB is necessarily a CECB.

In such games, any deviation leads to finite paths and, hence, in any deviation there exists a node at which the action changes, but the action at each node reachable from this node stays the same. The deviation in which only the action at this node changes would be one-step, credible, and profitable.

The one-step deviation property no longer holds in infinite games or games with cycles. The reason is that, under a (C)ECB, a deviation is profitable only if it improves the payoffs of the deviators at every node at which an action changes. Hence, in infinite games, for a deviation to be profitable, a coalition might need to promise to change its actions at infinitely many places. An example is available in Section 4.1.

3.3 *Computation in finite games and backward induction*

3.3.1 *Finite acyclic games* A nice implication of the one-step deviation property is that both ECB and CECB can be found through backward induction in finite acyclic games. Let Γ be a finite acyclic game. For any $z \in Z$, let $\Gamma(z)$ denote the subgame starting at z and let $l(\Gamma(z))$ be the length of this game (the length of the longest terminal path).

Backward induction works as follows. For any z for which $l(\Gamma(z)) = 1$, find a (C)ECB for $\Gamma(z)$ and let the path of play at z be denoted by $p(z)$. Suppose the (C)ECBs of every subgame of length $l < k$ are found and take any z with $l(\Gamma(z)) = k$. Let $\hat{\Gamma}(z)$ denote the reduced basic game in which the set of nodes is only those nodes that are directly reachable from z and the utility of taking an action (z, z', S) is given by the utility of $\{(z, z', S), p(z')\}$. Find the (C)ECB of this reduced basic game and let $p(z)$ denote the resulting path of play at z. Continue in this way until all of the nodes are exhausted.

The backward induction algorithm above finds all of the (C)ECBs of any finite acyclic game.

3.3.2 *Finite games with cycles* Even when the game has cycles, one can use backward induction to a *certain extent*. For instance, if a node z' is reachable from node z and if $\Gamma(z')$ is a proper subset of $\Gamma(z)$, then by the same argument as above, without loss of generality, one can first find the (C)ECB of $\Gamma(z')$ and then move up.[9]

Given this, the question is the computability of the (C)ECB in basic games with cycles. The question is particularly relevant for the CECB, as its definition requires us to check deviations from deviations. Chwe (1994) has already provided an algorithm to compute the largest consistent set, so we can easily adapt the algorithm to work for the CECB.

Assume that Γ is a basic game and let Φ be the set of coalitional behaviors. Define the function $f : 2^{\Phi} \to 2^{\Phi}$, where $f(X) = \{\phi \in \Phi \mid \text{for all } S \subseteq N \text{ and } \phi' \in \Phi \text{ such that } \phi' >_D^S \phi, \exists \phi^* \in X \text{ with } \phi^* \rhd \phi' \text{ and } \phi^* \not>_S \phi\}$.

Note that a set of coalitional behaviors X is credible if and only if $f(X) = X$. Furthermore, f is isotonic, i.e., $X \subseteq Y \implies f(X) \subseteq f(Y)$. We iteratively apply the function

[9]The example in Figure 2 clearly demonstrates this process: We can start solving this game from the "final" proper subgame of the game, which is the cycle depicted in Figure 2(c). After we find the CECB of this subgame, we can plug the utility of the CECB into node c, which connects this subgame to the game itself. Once we do this, we end up with the game in Figure 2(b). This is a basic game, which means that we cannot reduce it any further. Hence, we solve this game and then we are done.

f starting from Φ. Since f is isotonic, we have that $\Phi \supseteq f(\Phi) \supseteq f(f(\Phi))\ldots$. Since Φ is finite, there exists j such that $f^j(Z) = f^{j+1}(Z)$, implying that $f^j(Z)$ is a credible set. Clearly there cannot be a set K strictly containing $f^j(Z)$ that is also credible. But then by Lemma 1, $f^j(Z)$ is the LCRS. Hence a coalitional behavior ϕ is a CECB if and only if $\phi \in f^j(Z)$.

3.4 *A non-cooperative game*

A stationary subgame perfect equilibrium of a simple bargaining game defined over an extended coalitional game can be associated with the (C)ECBs of the underlying game, which provides further evidence that (C)ECB embodies the idea of maximality.

The extensive form game will be defined for extended coalitional games in a tree structure, but it is easy to see that the results can be generalized to any acyclic extended coalitional game. Furthermore, we assume that the extended coalitional game satisfies the following property: if $(z, z, \varnothing) \in A_z$, then $A_z = (z, z, \varnothing)$, i.e., if no-action is available at node z, then node z is a terminal node.

We start with the description of the extensive form. An order of players for each node is given. The game starts at the root z_0.

Step 1. The first player in the exogenously given order becomes the proposer.

Step 2. The proposer proposes an action available at the node.

Step 3. Each player active at that action sequentially accepts or rejects the proposal according to the given order.

Step 4. If everyone accepts, the action is taken. The game moves on to the corresponding node. Go back to Step 1 if it is not a terminal node. If it is terminal, the payoffs are realized.

Step 5. If someone rejects, then he becomes the new proposer. Go back to Step 2.

If the negotiations come to an end, i.e., if a terminal node is reached, then each player gets the utility corresponding to the utility of the terminal path the game visits. Ongoing negotiations provide the worst utility.

We say that a game is strongly acyclic if the game is acyclic and the dominance relation $>_D$ defined on Φ is acyclic. The first result shows that any ECB can be supported as a stationary subgame perfect equilibrium in finite strongly acyclic games. Proposition 3 provides further evidence that the predictions of ECB are indeed sensible, where sensibility is checked against the stationary subgame perfect equilibrium of a simple bargaining game.

Proposition 3. *In finite strongly acyclic games, any ECB ϕ can be supported as a stationary subgame perfect equilibrium outcome for some order of players.*

Both of the solution concepts can be found through backward induction, so one can restrict attention to trees of length 1. Suppose (z, x, S) is an ECB for such a tree. Suppose the order of players is such that someone in S comes first. Consider the stationary (incomplete) strategy in which everybody in S offers (z, x, S) and accepts this offer. Note that for any action $(z, y, T) \neq (z, x, S)$, either $(z, x, S) \succeq_S (z, y, T)$ or there exists $i \in T$ with $(z, x, S) \succeq_i (z, y, T)$. In the former case, everyone in S would like to prevent the action; in the latter case, there exists $i \in T$, who would be willing to reject the action for the sake of (z, x, S). This gives a rough idea as to how the given incomplete strategies can be completed in a way that supports (z, x, S) as the stationary subgame perfect equilibrium outcome, where strong acyclicity is needed to complete the described strategy. See the Appendix for the details of the proof.

One might wonder if there exists a condition that would relate the stationary subgame perfect equilibrium to the (C)ECB in a stronger way. It turns out that indeed there is. The condition is somewhat strong, but it provides us with full implementation of both ECB and CECB.

For a coalitional behavior ϕ, let $\mathcal{I}(\phi)$ denote the set of individuals that become active at some point according to ϕ, i.e., $\mathcal{I}(\phi) = \{i \in N \mid \text{for some } z \in Z, i \in \mathcal{I}(\phi(z))\}$, where $\mathcal{I}(\phi(z))$ is the initiator of $\phi(z)$.

Definition 12. A coalition S is *potent* if there exists a coalitional behavior ϕ^* such that

- $\mathcal{I}(\phi^*) = S$ and $\mathcal{I}(\phi) \cap S \neq \varnothing$ for all $\phi \neq \phi^*$
- for any $\phi \in \Phi$, either $\phi^* \succ_S \phi$ or $\phi^* \sim_S \phi$.[10]

The potent coalition has the ability to impose their most preferred coalitional behavior. Furthermore, for any coalitional behavior, the potent coalition has a member who is needed to impose that coalitional behavior.

Proposition 4. *Suppose* Γ *is a finite acyclic game that contains a potent coalition. Then* ϕ *is a stationary subgame perfect equilibrium outcome if and only if* ϕ *is a CECB if and only if* ϕ *is an ECB.*

The proposition states that the existence of a potent coalition makes certain predictions highly salient and that all three solution concepts predict these coalitional behaviors.

Since the predictions of all three solution concepts can be found recursively, Proposition 4 can also be applied recursively. Namely, one can show that when applying backward induction, each reduced game has a potent coalition. Then Proposition 4 still applies even if the whole game does not have a potent coalition. That is, in such games we still have equivalence between stationary subgame perfect equilibrium, ECB, and CECB. A good example is the domain of extensive form games of perfect information. In this domain, in general, no potent coalition would exist. But it is easy to see that in the process of backward induction, each reduced game definitely contains a potent coalition. Hence, the theorem applies on this domain.

[10]Remember that $\phi^* \succ_S \phi$ if $\phi^* \succ_i \phi$ for all $i \in S$. Similarly, $\phi^* \sim_S \phi$ if $\phi^* \sim_i \phi$ for all $i \in S$, where $\phi \succ_i \phi'$ ($\phi \sim_i \phi'$) if for every $z \in Z$ such that $\phi(z) \neq \phi'(z)$, $\sigma(z) \succ_i \sigma'(z)$ ($\sigma(z) \sim_i \sigma'(z)$).

4. Classes of games

In this section we look at applications to general classes of games that are being extensively used in the economics literature. We start by analyzing the (C)ECBs of *extensive form games,* which is followed by the *abstract game* and the *characteristic function games.*

4.1 *Extensive form games*

It is easy to see that an extensive form game of perfect information is an extended coalitional game.

Definition 13 (Extensive form games of perfect information). An extended coalitional game Γ is an extensive form game if the graph of Γ is a tree and for all $z \in Z$, the following statements hold.

- Either $A_z = \{(z, z, \varnothing)\}$ (i.e., z is a terminal node) or for all $(z, z', S) \in A_z$, we have $S = \{i\}$ for the same $i \in N$ (i.e., only one individual is active at each node).
- For $p', p'' \in \mathcal{P}_z$ and $i \in N$, $p' \succeq_i p''$ if and only if $(p, p') \succeq_i (p, p'')$ for any path p that ends at z.

The first condition basically states that a single individual is active at each node. The second condition is the usual consistency condition on the payoffs that is inherent in extensive form games: namely, if p' and p'' are two terminal paths that start at z, then an individual prefers p' to p'' if and only if she prefers (p, p') to (p, p'') for any path p ending at z.

In finite extensive form games, since both subgame perfect equilibrium and (C)ECB satisfy the one-step deviation property, it is easy to show that subgame perfect equilibrium and (C)ECB are equivalent. However, the equivalence breaks down in infinite games. But one can easily establish that every ECB is going to be a subgame perfect equilibrium.

Proposition 5. *In any finite extensive form game, ϕ is an ECB if and only if ϕ is a CECB if and only if ϕ is a subgame perfect equilibrium. Furthermore, if ϕ is an ECB for any extensive form game, then ϕ is a subgame perfect equilibrium.*

The reason why equivalence breaks down in infinite games is that unlike subgame perfect equilibrium, the one-step deviation property of (C)ECB is not inherited in infinite horizon games that are continuous at infinity, such as games with discounting. In games that are continuous at infinity, under the subgame perfect equilibrium, any profitable infinite deviation can be replaced with a profitable finite deviation. This is done by truncating the deviation after some period T. As the payoffs get less and less important in future periods the resulting deviation is still profitable. But under a (C)ECB, this may not be the case, because whatever T at which we truncate the devia-

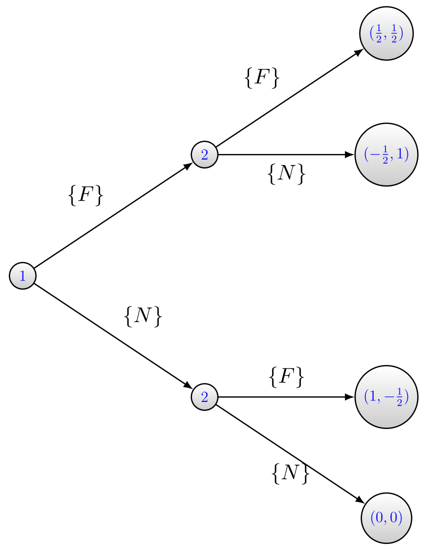

Figure 3. Favor exchange.

tion, (C)ECB also imposes the restriction that at period $T-1$, the deviation should be profitable.

For instance, consider the *favor exchange* game in Figure 3. Here, player 1 can do a favor for player 2 or not. Following player 1's decision, player 2 decides whether to do a favor for player 1 or not. A player who receives a favor gets a utility of 1 and a player who provides a favor incurs a cost of $\frac{1}{2}$.

There is a unique subgame perfect equilibrium of this game, which specifies that each player chooses to provide no favor at each node of the game. By Proposition 5, this is also the unique (C)ECB of the game.

Suppose that the players repeatedly play the favor exchange game. At each period they get the payoffs of the favor exchange game played in that period and the payoffs are discounted with a common discount factor $\delta \in (0, 1)$.

There is still a subgame perfect equilibrium of this game in which each player chooses $\{N\}$ in each period. But this subgame perfect equilibrium is not a (C)ECB for δ big enough. This is because there is a profitable deviation by $S = \{1, 2\}$ in which the players change to playing $\{F\}$ as long as no $\{N\}$ is chosen. The deviation is profitable since at every node at which S changes its action, both of the players are better off. Furthermore, the deviation is credible, because if any player chooses to deviate from the newly prescribed action, then he will be punished by the perpetual play of $\{N\}$.

This example also shows that unlike subgame perfect equilibrium, under (C)ECB, the one-step deviation property no longer holds in infinite horizon extensive form games with discounting. The above deviation cannot be replaced by a finite deviation, because in any finite deviation, the last player to deviate would have no incentive to keep his promise.

4.2 *The abstract game*

An abstract game is defined as $\Gamma = \{N, Z, \{v_i\}_{i\in N}, \{\overset{S}{\rightarrow}\}_{S\subseteq N, S\neq\varnothing}\}$ (see, for example, Chwe 1994, Rosenthal 1972, and Xue 1998), where N is the set of players, Z is the set of states, $\{v_i\}$ is player i's utility function defined on the set of states, and $\{\overset{S}{\rightarrow}\}_{S\subseteq N, S\neq\varnothing}$ are effectiveness relations defined on Z. The effectiveness relation $\{\overset{S}{\rightarrow}\}$ describes what coalition S can do at every state, i.e., $a \overset{S}{\rightarrow} b$ for $a, b \in Z$ if and only if when a is the status quo coalition, S can exchange state a with state b.

All of the ingredients of an abstract game, except for the utilities, admit an obvious translation to the extended coalitional game. The set of players N and the set of nodes Z at the same. The set of actions available at any node is given by the effectiveness relation, but also note that in an abstract game, it is possible for no-action to be taken at every node. Hence, for each $z \in Z$, A_z is defined as $(z, x, S) \in A_z$ if and only if (($x = z$ and $S = \varnothing$) or ($z \overset{S}{\rightarrow} x$ for $S \neq \varnothing$)).

By changing the way we define the utilities over the paths, (C)ECB is able to mimic the way different solution concepts analyze the abstract game. Main approaches can be divided into two.

The static approach Some solution concepts such as largest consistent set (LCS), farsighted stable set (FSS) (see Chwe 1994), optimistic stable standard of behavior (OSSB), conservative stable standard of behavior (CSSB) (see Xue 1998), REFS, and SREFS (see Dutta and Vohra 2017) assume that players care only about the final outcome the path leads to. This might be easily represented with the following assumption on preferences: For any $p \in \mathcal{P}$, let $\mathcal{T}(p)$ denote the node in which the path terminates if p is finite; otherwise let $\mathcal{T}(p) = \varnothing$. For all $i \in N$ and $p \in \mathcal{P}$, we set $u_i(p) = v_i(\mathcal{T}(p))$ if p is finite; otherwise $u_i(p) = -\infty$. I call the (C)ECB under this assumption the *static (C)ECB*.[11]

The dynamic approach Some solution concepts such as the EPCF (see Konishi and Ray 2003 and Ray and Vohra 2014) assume that there is a discount factor δ and players discount the utilities among a path. This might be easily represented with the following assumption on preferences: For $p = \{z_k\}_{k=1,\dots,K} \in \mathcal{P}$, let $u_i(p) = \sum_{k=0,\dots,K} \delta^k v_i(z_k) + \sum_{k=K+1,\dots} \delta^k v_i(z_K)$. I call the (C)ECB under this assumption the *dynamic (C)ECB*.

Early approaches to study farsighted coalition formation suffered from what Ray and Vohra (2014) call the problem of maximality. This refers to the observation that instead of considering the best course of action, players form extreme expectations based on optimism (such as FSS and OSSB) or pessimism (such as LCS and CSSB).

The example in Figure 4 demonstrates this for the LCS. In this example, the LCS is $\{a, c, d\}$. According to LCS, a is stable because under the assumptions of the LCS, players are pessimistic and player 1 is afraid that player 2 will move to state c following state b. It is easy to see that (C)ECB makes the "correct" prediction in this game.[12]

[11] Assigning the lowest utility to infinite paths causes no loss in generality, as the solution concepts in the static approach never prescribe cycles or infinite paths.

[12] The unique (C)ECB ϕ specifies $\phi(a) = (a, b, \{1\})$ and $\phi(b) = (b, d, \{1, 2\})$.

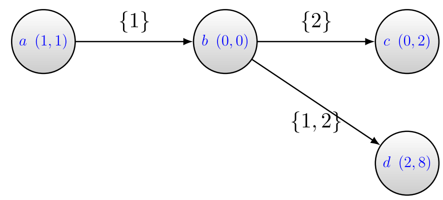

FIGURE 4. LCS differs from (C)ECB.

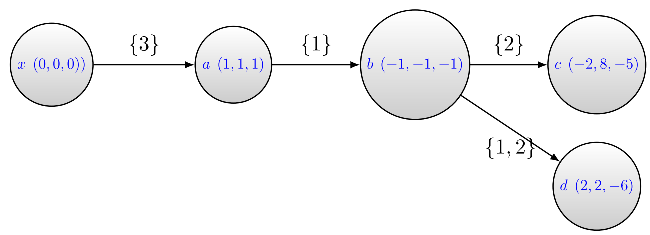

FIGURE 5. FSS differs from (C)ECB.

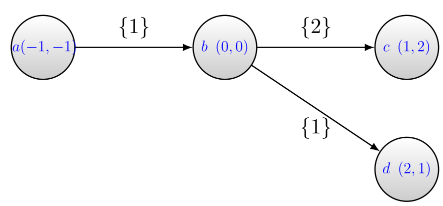

FIGURE 6. OSSB/CSSB differ from (C)ECB (example taken from Herings et al. 2004).

Similarly, the example in Figure 5 demonstrates the problem for the FSS. The reasonable prediction and the unique (C)ECB ϕ of this game specify $\phi(b) = (b, c, \{2\})$, $\phi(a) = (a, a, \varnothing)$, and $\phi(x) = (x, a, \{3\})$. However, under the FSS, a is not stable, because 1 optimistically and unreasonably believes that 2 will move to d with him, but this in turn makes x stable under the unique FSS. Hence, FSS $= \{x, c, d\}$.

Xue (1998) proposed to deal with the issues related to the FSS and the LCS by endowing individuals with perfect foresight in the sense that under Xue's (1998) solution concepts, individuals not only consider the final outcomes, but also consider how these outcomes are reached. For this reason, Xue (1998) looks at the stability of a set of paths instead of the stability of a set of outcomes while defining his solution concepts OSSB and CSSB.

One problem with OSSB and CSSB is that both of these solution concepts might fail to make some obvious predictions, clearly violating the notion of maximality. Nice and simple examples of these problems are provided in Herings et al. (2004). In all of these examples it is easy to see that (C)ECB makes the right prediction. Here I replicate only one of the examples provided by Herings et al. (2004).

Consider the game in Figure 6 taken from Herings et al. (2004). The authors note that in this example, under the unique OSSB/CSSB state, a is stable. This should not be

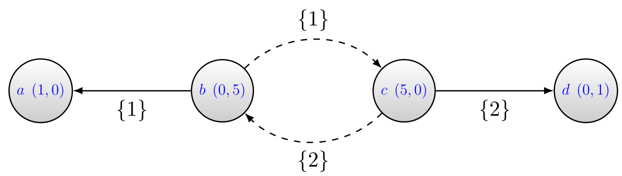

FIGURE 7. Profitable deviation to a cycle.

the case, because if player 1 moves to b, then she will certainly be better off. The reason OSSB/CSSB fails is that at node b, both the OSSB and the CSSB are empty and this leads both concepts to fail to make the right prediction at the node preceding node b. Finally, it is easy to see that (C)ECB makes the correct prediction in this game. Indeed, there are two (C)ECBs, ϕ_1 and ϕ_2, where $\phi_1(a) = \phi_2(a) = (a, b, \{1\})$, $\phi_1(b) = (b, d, \{1\})$, and $\phi_2(b) = (b, c, \{2\})$.

Nevertheless, some recent solution concepts such as the REFS, SREFS, and EPCF have managed to solve these issues by considering consistent expectations just like (C)ECB does. One important way in which (C)ECB differs from these solution concepts on the domain of the abstract game is that all these solution concepts restrict attention to one-step deviations, whereas (C)ECB allows for any arbitrary deviation.

This does not seem to be a problem in static solution concepts such as the REFS and SREFS. However, the same cannot be said of the dynamic solution concepts such as the EPCF.

To see this, consider the example in Figure 7. Assume that the utilities are defined as discounted utilities over paths. At node b, player 1 moves to a and settles for a payoff of 1. Similarly, player 2 moves from c to d and settles for a payoff of 1. There is no profitable one-step deviation from this coalitional behavior for δ big enough. But there is a profitable deviation by the coalition $\{1, 2\}$ to the cycle. Hence, for δ big enough, the depicted coalitional behavior is not a (C)ECB, but it is an EPCF, although the profitable deviation to the cycle would increase the payoff of every individual.

As has been demonstrated, one advantage of (C)ECB is its versatility. By changing the way we define the utilities over the paths, (C)ECB is able to mimic the way different solution concepts analyze the abstract game.

Indeed, one can easily show that ECB is related to a solution concept in both the static approach (SREFS) and the dynamic approach (EPCF). When one also includes the relationship between (C)ECB and SPE, one obtains a solution concept that is able to bridge the non-cooperative approach with the static and the dynamic approaches to foresight. Now we show this formally.

DEFINITION 14 (Indirect dominance). The relationship $x \in Z$ indirectly dominates $y \in Z (x \gg y)$ if there exists $x_0, x_1, \dots, x_n \in Z$ and $S_0, S_1, S_2, \dots, S_{n-1}$ such that $x_0 = y$, $x_n = x$, $(x_i, x_{i+1}, S_i) \in A_{x_i}$, and $v_j(x_n) > v_j(x_i)$ for all $j \in S_i$, for all $i = 0, \dots, n-1$.

The solution concepts in the static approach are defined using the indirect dominance relation, which uses strict preference, whereas under the (C)ECB, coalitions might take some actions even if they are indifferent. We always see a difference between these concepts and the (C)ECB due to the different ways with which they treat indifference. I do not view this difference as essential. Therefore, not to bump on it again and again, from now on I assume that the environment satisfies *no indifference*.

Definition 15 (No indifference). An abstract game Γ satisfies *no indifference* if for all $i \in N$ and $z, z' \in Z$, where $z \neq z'$, we have $v_i(z) \neq v_i(z')$.

Using coalitional behavior rather than expectations, we can restate the stability concept used in SREFS in the language of our framework. Let $S(\phi)$ denote the set of stable nodes in coalitional behavior ϕ, i.e., $S(\phi) = \{z \in Z \mid \phi(z) = (z, z, \varnothing)\}$.

Definition 16 (SREFS (Dutta and Vohra 2017)). A set $V \subseteq Z$ is an SREFS if there exists an acyclic coalitional behavior ϕ such that $S(\phi) = V$ and the following statements hold.

(**IS**) If $x \in V$, then there does not exist $y \in Z$ and $S \subseteq N$ such that $(x, y, S) \in A_x$ and $v_i(\mathcal{T}(\sigma(y))) > v_i(x)$ for all $i \in S$.

(**ES**) If $x \notin V$, then $\sigma(x)$ is an indirect dominance path.

(**M**) If $x \notin V$ and if T is the initiator at x, then there does not exist $y \in Z$ and $F \subseteq N$ with $T \cap F \neq \varnothing$ and $(x, y, F) \in A_x$ such that $v_i(\mathcal{T}(\sigma(y))) > v_i(\mathcal{T}(\sigma(x)))$ for all $i \in F$.

The first and second conditions are interpreted as internal and external stability conditions with respect to the expectation, whereas the third condition requires optimality of the move at any node x, where optimality is conditioned on a one-step deviation.

It turns out that under the weak assumption that actions are monotonic, in the sense that whenever a coalition S is able to take a certain action, then any coalition T containing S can also take this action, ECB is equivalent to SREFS.

Definition 17 (Monotonicity of actions). We say that an extended coalitional game satisfies *monotonicity of actions* if whenever $(z, z', S) \in A_z$ for some $z \in Z$, we also have that $(z, z', T) \in A_z$ for all $T \supseteq S$.

Proposition 6. *Let Γ be an abstract game that satisfies no indifference and monotonicity of actions. Then the following statements hold.*

- *If V is an SREFS and ϕ is the coalitional behavior that supports it, then ϕ is a static ECB and, hence, is a static CECB.*
- *If ϕ is a static ECB, then $S(\phi)$ is an SREFS supported by ϕ.*

This establishes that (C)ECB (when defined as a static concept) is closely related to a solution concept in the static approach and in turn it also means that static (C)ECB's predictions also satisfy the internal and external stability properties imposed by this solution concept. On the side of SREFS, this result shows that the conditions of internal and external stability with the one-step deviation condition rule out every possible profitable deviation, even those that involve multiple actions.

EPCF also uses consistent expectations just like SREFS and (C)ECB. The difference between EPCF and SREFS (apart from one being a static and the other being a dynamic concept) is that the former is directly defined as an expectation (a coalitional behavior) that is immune to certain deviations. Using coalitional behavior, we can restate the definition of EPCF under *no indifference* in the language of our framework.[13]

Definition 18 (EPCF). A (deterministic) EPCF[14] is a coalitional behavior ϕ such that for all $x \in Z$, the following statements hold, where for any $z, z' \in Z$, $\sigma(z) \succ_i \sigma'(z)$ if and only if the discounted utility of the former is greater.

- If $\phi(x) = (x, y, S)$, where $y \neq x$, then $\sigma(y) \succ_S \sigma(x)$ and there does not exist z with $(x, z, S) \in A_x$ and $\sigma(z) \succ_S \sigma(y)$.
- If x is such that there exists $y \in Z$ and $S \subseteq N$ with $(x, y, S) \in A_x$ and $\sigma(y) \succ_S \sigma(x)$, then $\phi(x) \neq (x, x, \varnothing)$.

The definition is similar to the definition of dynamic (C)ECB with the major difference being that this definition does not consider deviations that involve multiple actions.[15] With the dynamic assumption on the utilities, dynamic (C)ECB does not satisfy the one-step deviation property. This means that the restriction to one-step deviations in the above definition is with loss of generality. In particular, there might exist a deviation to a cycle that might make the deviators better off (see Figure 7), which also implies that there will be EPCFs that are not ECBs (again see Figure 7 for an example).

Finally, it is easy to establish that every dynamic ECB is an EPCF.

Proposition 7. *If ϕ is a dynamic ECB, then it is an EPCF, but an EPCF may not be a dynamic ECB.*

When one adds the relationship of ECB to subgame perfect equilibrium to these, one gets a solution concept that is directly or indirectly related to a wide area of the literature on farsighted coalition formation; see Figure 8.

Hence, I believe that one major contribution of this paper is to show that the developments in the literature have made it possible to achieve some level of unification between both the non-cooperative and the cooperative strands of the literature and within the cooperative approach to farsighted coalition formation.

[13]The assumption of no indifference is needed here for an altogether different reason. In particular in the original definition of EPCF, Konishi and Ray allow a coalition to move at state x even if it is indifferent between moving or staying at x. But they do not allow this if there exists a coalition that can move at x and that would strictly improve by taking an action at x, whereas under a (C)ECB, this is also allowed.

[14]Konishi and Ray's (2003) EPCF can also be stochastic, as mixing is not included in the definition of an ECB. I restrict attention to deterministic EPCFs.

[15]Another difference is that the definition of the deviations is, in general, weaker under EPCF.

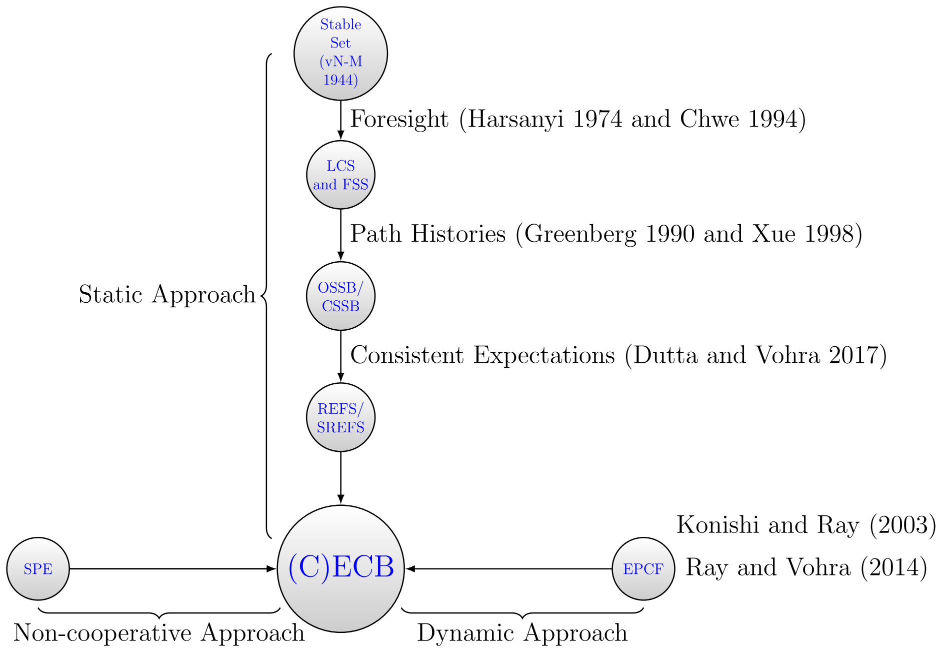

Figure 8. An overview.

4.2.1 *Other approaches on the abstract game* Some solution concepts such as those studied in Herings et al. (2004) and Granot and Hanany (2016) propose to define solution concepts on the abstract game by directly defining a non-cooperative game from the abstract game. Specifically, Granot and Hanany (2016) model the evolution of play resulting from coalitional deviations as an infinite extensive form game and define their solution concept—*subgame perfect consistent set* (SPCS)—as all the states that can be supported by a subgame perfect equilibrium of the extensive form, refined to additionally satisfy internal and external stability in line with the stable set. They show the existence of the solution concept in finite acyclic games and show that surprisingly the SPCS leads to efficiency in any normal-form game having a pure strategy Nash equilibrium.

Herings et al. (2004) take the abstract game as the primitive and define a multistage game associated with the abstract game. Then they define an appropriate notion of rationalizability on this multistage game. In an earlier version (see Herings et al. 2000), they define the notion of social rationalizability directly on the abstract game. The solution aims to be weak and identifies the coalitions that are likely to form and the outcomes that are likely to occur if the individuals are rational and they are endowed with a hierarchy of hypothesis, which is common knowledge at the status quo. They show that the set of socially rationalizable outcomes is nonempty and it satisfies a coalitional rationality property: that when the outcomes are Pareto ranked, it always chooses the Pareto dominant outcome.

The approaches taken in Granot and Hanany (2016) and Herings et al. (2000, 2004) is fruitful, but also fundamentally different than the approach I am taking with the (C)ECB and the other approaches discussed so far, where the solution concept is directly defined over the domain with no use of individual beliefs or strategies. That is, all of the solution concepts defined up to now, including the (C)ECB, directly make their predictions by considering coalitional actions/deviations, whereas Granot and Hanany (2016) appeal directly to a non-coooperative solution concept and Herings et al. (2000, 2004) appeal to individual beliefs to define their solution concept. In this sense, (C)ECB and the solution concepts of the dynamic and static approaches are firmly situated as "cooperative solution concepts," where cooperation is not explained through the help of a non-cooperative game or individual beliefs/strategies, whereas in Herings et al. (2004) and Granot and Hanany (2016), this is not the case.

These approaches should be seen as complementary and each would have its own advantages. Nevertheless, here I would like to mention some properties that I regard as advantages of (C)ECB over the solution concepts in Granot and Hanany (2016) and Herings et al. (2004).

First, appealing to the non-cooperative approach increases the complexity of the solution concept to a great extent. This is especially true when we compare these concepts to ECB, but also true to an extent for CECB. Consider the following example taken from Herings et al. (2000, 2004).

Example 1. We have $N = \{1, 2, \ldots, n\}$ and $Z = \{x_0, x_1, \ldots, x_K\}$, and for all $i \in N$ and $k \neq 0, K$, we have $u_i(x_K) > u_i(x_k) > u_i(x_0)$; x_0 is assumed to be the status quo. From x_0, N can move to any x_i, where $i = 1, 2, \ldots, K$. It is assumed that the payoff of a path is the payoff of the final outcome on the path.

The example above is a simple example in which N chooses between outcomes, one of which strictly Pareto dominates the others. It is *immediate* that the unique (C)ECB in this game specifies that N is going to move to the Pareto optimal outcome x_K. Social rationalizablity of Herings et al. (2000, 2004) also makes the same prediction, but it is much harder to show; it takes the authors several pages to come to this conclusion. The same critique, although in a different way, also applies to Granot and Hanany (2016). The fact that their solution concept requires the formation of an *infinite extensive form game* to solve the above simple example implies a considerable complication on the side of their solution concept.

Second, (C)ECB and its domain allows us much more flexibility than these solution concepts. Indeed, both social rationalizability and SPCS are *static* solution concepts in the sense that payoffs are assumed to be realized only when a stable state is reached, whereas (C)ECB allows for much richer preferences over the paths.

This flexibility allows (C)ECB to unify the different approaches to farsighted coalition formation. The approaches unified by (C)ECB range from the noncooperative approach to the static and dynamic approaches on the abstract game. This is the most clear contribution of (C)ECB over these other solution concepts.

4.3 *Characteristic function games*

A characteristic function game is a pair (N, V), where N is the finite set of players and for each coalition $S \subseteq N$, $V(S) \subseteq \mathbb{R}^S$ denotes the set of payoff vectors achievable by coalition S. A coalition structure P is a partition of N. A state z is a pair (x, P), where P is a coalition structure and x is an allocation that satisfies $x_S \in V(S)$ for all $S \in P$. Let Z denote the set of all states. The core is probably the most well known solution concept defined for characteristic function games. It is the set of states that no coalition can improve upon.

Definition 19. The core of the game (N, V) is defined as

$$C(N, V) = \{(x, P) \in Z \mid \text{there does not exist } S \subseteq N \text{ and } y_S \in V(S) \text{ such that } y_S > x_S\}.$$ [16]

I am going to study characteristic function games in line with the farsighted coalition formation literature, which assumes that a deviation by a coalition might be followed by any deviation by any other coalition and the players are farsighted in the sense that they care about the final allocation the negotiations lead to.

Basically, I translate the abstract game used in Ray and Vohra (2015) to an extended coalitional game. For any $(x, P) \in Z$, Ray and Vohra (2015) require the set of actions to satisfy the three conditions below. (Also see Konishi and Ray 2003 and Kóczy and Lauwers 2004, who use similar conditions.)

(i) We have $((x, P), (x, P), \varnothing) \in A_{(x,P)}$.

(ii) If $((x, P), (y, P'), S) \in A_{(x,P)}$, then $y_S \in v(S)$ and if $T \in P$ is such that $S \cap T = \varnothing$, then $T \in P'$ and $x_T = y_T$.

(iii) For all $(x, P) \in Z$, $T \subseteq N$ and $z_T \in V(T)$ such that either $z_T \neq x_T$ or $T \notin P$, there exists $((x, P), (y, P'), T) \in A_{(x,P)}$ such that $T \in P'$ and $y_T = z_T$.

The first condition states that it is possible to stay in every state. The second condition requires that when a coalition deviates from an outcome, it has to get something feasible for itself and it cannot dictate the payoffs and structures of the coalitions that are unrelated to it. Finally, the third condition states that if a payoff z_T is feasible for a coalition T ($z_T \in V(T)$), then T should be able to get z_T or if a coalition T has not formed, then T should be able to form.

Hence, an extended coalitional game that corresponds to a characteristic function game is $\Gamma = \{N, Z, \{A_z\}_{z \in Z}, \{u_i\}_{i \in N}\}$, where N is the set of players, Z corresponds to all states, $\{A_z\}_{z \in Z}$ is any set of actions that satisfy the restrictions above, and the utilities correspond to the static approach, i.e., for all $i \in N$ and $p \in \mathcal{P}$, we set $u_i(p) = v_i(\mathcal{T}(p))$ if p is finite; otherwise $u_i(p) = -\infty$, where $v_i((x, P)) = x_i$.

CECB is not a very satisfactory concept for characteristic function games for a particular set of reasons: (a) Characteristic function games are basic games; hence, backward induction cannot be applied even at the basic game level. This makes the solution concept too permissive on characteristic function games. (b) Because of the multitude of

[16] We have $y_S > x_S$ if $y_i > x_i$ for all $i \in S$.

coalitional behaviors possible in such games, it is very hard to find the set of CECBs. For these reasons, I only analyze the ECB in this class.

The proposition below establishes that if, under an ECB, the path of play from every node terminates at the same state, then this state is in the core. Furthermore, if a state is in the core, then there exists an ECB such that the path of play from every node terminates at this core state. That is, ECBs with a single prediction completely characterize the core.

Proposition 8.

- *If* $(x^*, P^*) \in C(N, V)$*, then there exists an ECB* ϕ *such that* $\mathcal{T}(\sigma(x, P)) = (x^*, P^*)$ *for every* $(x, P) \in Z$*.*
- *If* ϕ *is an ECB such that* $\mathcal{T}(\sigma(x, P)) = (x^*, P^*)$ *for every* $(x, P) \in Z$*, then* $(x^*, P^*) \in C(N, V)$*.*

The proposition establishes that a farsighted solution concept completely characterizes a well known myopic concept: the core. The result also complements other results in the literature that are close in spirit to this result. Other results that show that the core incorporates foresight include Ray (1989), Diamantoudi and Xue (2003), Konishi and Ray (2003), Mauleon et al. (2011), and Ray and Vohra (2015).[17] Ray (1989) shows that the core is immune to nested objections. The results in Diamantoudi and Xue (2003), Mauleon et al. (2011), and Ray and Vohra (2015) concern the FSS, whereas Konishi and Ray's (2003) result is related to the EPCF.

Proposition 8 implies that if a characteristic function game has an empty core, then we cannot find an ECB for that game with a single stable outcome. But this does not mean that an ECB does not exist in such a game. An example is available from the author upon request.

5. An application: Network formation

There is a tension between efficiency and stability in myopic models of network formation (see Jackson 2010 and Jackson and Wolinksy 1996). Furthermore, Dutta et al. (2005) and Herings et al. (2009) have shown that this tension continues in farsighted solution concepts.

In this section, we see that (C)ECB helps to alleviate this tension to a great extent. In particular, Proposition 9 shows that whatever the value of the network is, as long as it satisfies a mild assumption, there is a reasonable way to allocate this value to each player so that every efficient network can be supported by both the dynamic ECB and the dynamic CECB if the players are sufficiently patient.[18] This proposition is almost completely the opposite of Theorem 1 in Jackson and Wolinsky (1996), Theorem 2 in

[17] Green (1974), Feldman (1974), Kóczy and Lauwers (2004), and Sengupta and Sengupta (1996) show how myopic objections lead to the core in specific environments. Although these papers are also related, they are fundamentally different from the current paper, since players are not assumed to be farsighted.

[18] This result would still hold with the static (C)ECB.

Dutta et al. (2005), and Theorem 6 in Herings et al. (2009), which show the tension between stability and efficiency in a myopic stability concept in an approach based on the EPCF and in FSS, respectively.

But then, what allows (C)ECB to achieve this "possibility of efficiency" result that is different from the results in the literature? The idea is that a certain simple way to allocate the value of the network allows the players to punish any coalition that deviates from the efficient network, while rewarding the punishers with a payoff higher than what they would get under the efficient network. The key here is nonstationarity, which allows the players to punish the deviator. Hence, the actions assigned by the (C)ECB not only depend on the network structure, but also on the identity of the deviators.

Building upon this result, Proposition 10 provides a simple condition under which every prediction of (C)ECB is an efficient network structure if the players are sufficiently patient.

Network formation games can be seen as a generalized version of characteristic function games. From Section 4.3, we know that ECBs with a unique outcome completely characterize the core of a characteristic function game. Is there a corresponding result for network formation games?

It is not even immediate what core is in this context, because there might be widespread externalities across links. Nevertheless, we can show that (see Proposition 12) *static* ECBs with a unique prediction completely characterize the *pessimistic network core*, which assumes that when a coalition deviates, it can only form a network within itself.

This result is only partially true with the *dynamic* ECB. The reason is that the definition of the core ignores deviations to cycles, and in a dynamic ECB, it is quite possible that the players may prefer to cycle between different network structures (see Example 5).

The *bottom line* is that (C)ECB does more than unify the alternative approaches to foresight. This section demonstrates this through showing that in an application that is quite general and that has already been extensively studied, (C)ECB can make novel predictions, some of which go against the findings in the literature.

This is not to say that (C)ECB is better than the other approaches to foresight. Indeed, how can we claim this given that it is equivalent to certain prominent solution concepts under mild assumptions? But this section clearly shows that there is value added in the *flexibility* provided in (C)ECB beyond the unification aspect of the solution concept. This flexibility allows us to establish results that have been overlooked in the literature.

5.1 *Preliminaries*

Let $N = \{1, 2, \dots, n\}$ be a finite set of players. The complete graph on N is denoted by g^N. The set of all possible networks (undirected graphs) on N is denoted by $\mathcal{G} = \{g \mid g \subseteq g^N\}$. A component of a network g is a subset c of g such that no $i \in c$ is linked to $j \notin c$ and every distinct i and j in c are either directly or indirectly linked to each other. Note that this definition treats isolated players as components. The set of all components of g is denoted by $C(g)$.

The value of a network is given by a value function $v : \mathcal{G} \to \mathbb{R}$. The set of all value functions is denoted by V. A value function is *anonymous* if, for any permutation of the set of players π, $v(g^\pi) = v(g)$, where $g^\pi = \{\pi(i)\pi(j) \mid ij \in g\}$. A value function is *component additive* if $v(g) = \sum_{\{c \in C(g)\}} v(c)$ for all $g \in \mathcal{G}$. Throughout, the value function is assumed to be component additive.

An allocation rule is a function $Y : \mathcal{G} \times V \to \mathbb{R}^n$ such that $\sum_{i \in N} Y_i(g, v) = v(g)$. It describes how the value of the network is distributed among individuals. An allocation rule Y is anonymous if for any $v \in V$, $g \in \mathcal{G}$, and permutation of players π, $Y_{\pi(i)}(g^\pi, v^\pi) = Y_i(g, v)$, where v^π is defined as $v^\pi(g^\pi) = v(g)$. An allocation rule Y is component balanced if for any $v \in V$, $g \in \mathcal{G}$, and $c \in C(g)$, $\sum_{i \in c} Y_i(g, v) = v(c)$.

The componentwise egalitarian allocation rule, denoted by Y^{ce}, is defined as follows. For any $v \in V$, network g, and $i \in N$,

$$Y_i^{\text{ce}}(g, v) = \frac{v(c)}{|c|}, \quad \text{where } c \in C(g) \text{ is such that } i \in c.$$

Finally, we say that a network g is efficient if $v(g) \geq v(g')$ for all $g' \in \mathcal{G}$. Note that this is a strong definition of efficiency; in particular, it is stronger than Pareto efficiency.

5.2 *The extended coalitional game*

I model network formation as an infinite horizon tree, where the payoffs are realized dynamically. Hence, two important modeling decisions are made: (i) to study the game as an infinite horizon tree as opposed to a game with cycles and (ii) to analyze the dynamic (C)ECB as opposed to the static (C)ECB. On the latter, studying the dynamic (C)ECB as opposed to the static (C)ECB is without loss of generality in the context of my results. Remark 2 explains that all of the results of this section still hold if we consider the static (C)ECB instead of the dynamic (C)ECB.

On the former, modeling network formation as an infinite tree allows us to capture the richness of the process. For instance, modeling the game as a finite cyclic game would impose stationarity, whereas here we allow for the action specified at a particular network to depend on the set of actions that lead to that particular network structure. See Remark 3 for a discussion of this point.

Now I formalize the extended coalitional game. Each node z of the extended coalitional game is represented by a pair (g, h), where $g \in \mathcal{G}$ is the associated network and h denotes the history, i.e., the set of actions that lead to the node z from the root of the tree. For any histories h and h', let $h \cdot h'$ denote the concatenation of h with h'. The game starts at $z_0 = (g^\varnothing, \varnothing)$, where $g^\varnothing$ is the empty network and $h = \varnothing$ denotes the empty history.[19] At any node $z = (g, h)$, it is possible for no coalition to move, but also coalition S can move the game to $z' = (g', h \cdot (gg', S))$ if $g' \neq g$ and

(i) $ij \in g'$ and $ij \notin g$ implies $i, j \in S$

(ii) $ij \in g$ and $ij \notin g'$ implies $\{i, j\} \cap S \neq \varnothing$.

[19]This is without loss of generality. The results of this section would hold no matter where the game starts.

Hence, I allow for coalition-wise deviations. Most of the literature (including the farsighted literature on network formation) analyzes network formation with pairwise deviations. I focus on coalition-wise deviations, because doing so allows me to present the results in a more straightforward manner. With pairwise deviations we are able to get similar results, although the analysis is complicated. Remark 4 further elaborates on this point.

Finally, the payoffs are obtained in real time as in Dutta et al. (2005). That is, our focus in on the dynamic (C)ECB, where the payoff of a terminal path is the discounted utility of the networks visited along the path with a common discount factor δ.

5.3 *Efficiency and (C)ECB*

The literature has shown that there is a trade-off between efficiency and stability. Furthermore, simply considering farsighted solution concepts does not help with this trade-off to a great extent. In particular, there is no allocation rule that is component balanced and anonymous such that for each value function v, at least one efficient network is in the set of pairwise stable networks (see Jackson and Wolinsky 1996) and the farsighted stable set (see Herings et al. 2009).[20] Dutta et al. (2005) show the same result for a solution concept based on EPCF, assuming anonymity and limited transfers instead of anonymity and component balance, where an allocation rule permits limited transfers if $Y_i(g, v) \leq v(g)$ for all $i \in N$.

The proposition below shows that with (C)ECB, these results are no longer valid under a mild assumption. With the simple componentwise egalitarian rule, which is anonymous, is component balanced, and allows limited transfers, each efficient network can be supported as the prediction of a (C)ECB as long as the value function is anonymous. Note that g is a prediction of a (C)ECB ϕ if the path of play from the root of the game terminates at g, i.e., $\mathcal{T}(\sigma(z_0)) = g$.

Proposition 9. *There exists a rule Y that is anonymous, is component balanced, and allows limited transfers such that for each anonymous $v \in V$, every efficient network can be supported as the prediction of a (C)ECB for δ large enough. One such rule is Y^{ce}.*

The proof is lengthy, but with a simple basic idea behind it, which nicely demonstrates the workings of (C)ECB in this environment and how and why it makes different predictions. If we take the componentwise egalitarian rule, then for any efficient network g^* and for any player i, we can find a punishment network g_i that would decrease i's utility while increasing the utility of the players needed to sustain g_i. But then we can avoid a deviation by any player i through the threat of g_i. Similarly, we can avoid a deviation by any coalition S by singling out a player $i \in S$ as the would-be punished player in the case of a deviation by S.

[20]The result in Herings et al. (2009) is more general than this and applies to a solution called *pairwise farsightedly stable sets*. But they show that each farsighted stable set is also a *pairwise* farsightedly stable set.

The following example demonstrates the construction in the proof. The value function used in the example is the same as that used by Herings et al. (2009) to show the opposite result.

Example 2. Let $N = \{1, 2, 3\}$. Take the value function $v(12, 13, 23) = 9$, $v(12) = v(13) = v(23) = 8$, and $v(g) = 0$ for any other network.

We can easily find Y^{ce} for this value function: every player gets a payoff of 3 in the complete network, each connected player gets a payoff of 4 in a network with a single link, and every player gets a payoff of 0 in all remaining networks.

The efficient network is the complete network, but any two player coalition has an incentive to deviate to the network where the other player is isolated. The construction in the proof avoids such a deviation by the punishment that leaves one of the deviators as the isolated player indefinitely. Note that the punishment also rewards the other players by giving them the best payoff they can possibly get. Through this punishment, the efficient network can be supported indefinitely if the players are sufficiently patient.

The following example shows that v needs to be anonymous for the proposition to hold. When v is not anonymous, one might not be able to find punishment strategies to support the efficient network under Y^{ce}.

Example 3. Let $N = \{1, 2, 3\}$. Take the following value function: $v(12, 13, 23) = 9$, $v(12) = 8$, and $v(g) = 0$ for any other network. This v is not anonymous as $v(13) = v(23) = 0$, while $v(12) = 8$.

With this value function, the efficient network is the complete network. But under Y^{ce}, the unique CECB (so, ECB) specifies that the coalition $\{1, 2\}$ directly moves to the network $g = \{\{12\}\}$ and stays there indefinitely. The reason that the efficient network cannot be supported is that player 3 is a weak player and he cannot be a part of any punishment to the deviation by coalition $\{1, 2\}$.

Although the efficient network can always be supported with a (C)ECB, it is not the unique prediction. Indeed, we can see this in Example 2, where the efficient network is supported through the threat of a network that is not efficient; hence, the punishment network can also be supported as the prediction of a (C)ECB. So under what conditions can we find an allocation rule such that all of the predictions of a (C)ECB are efficient? The following proposition provides an answer to this question.

Let g^S denote the complete network on the set of players S. For any v and $S \subseteq N$, let $p(v, S) = \max_{g \subseteq g^S} \frac{v(g)}{|S|}$. We say that a value function v is *strictly top convex* if $p(v, N) > p(v, S)$ for all $S \neq N$. Strict top convexity is a simple modification of Jackson and van den Nouweland's (2005) *top convexity*, where the only difference is that the strict inequality is replaced with a weak inequality.

Proposition 10. *Suppose v is strictly top convex. Then there exists a discount factor δ^* and a rule Y that is anonymous and component balanced such that g is a prediction of a (C)ECB for every $\delta \in (\delta^*, 1)$ if and only if g is efficient.*

To prove this result, I use the rule Y^{ce}. We already showed that under this rule, any efficient network can be supported as the prediction of a (C)ECB. Hence, all we need to do is to show that each prediction is efficient. Note that when v is strictly top convex, under Y^{ce}, each efficient network also (strictly) maximizes the payoff of each player. Hence, if the efficient network is not predicted from some network g', then the players can deviate and directly impose an efficient network from g'. The deviation is certainly profitable, as the efficient network (strictly) maximizes the payoff of each player. For the same reason, the deviation is also going to be credible.

Jackson and van den Nouweland (2005) show that the set of strongly efficient networks is the set of strongly stable networks under the componentwise egalitarian allocation rule if v is top convex. In addition, Grandjean et al. (2011) show that the set of strongly efficient networks is also the unique farsightedly stable set under the componentwise egalitarian allocation rule if v is top convex. Together with Proposition 10, we see that when v is strictly top convex, the set of strongly stable networks, the farsighted stable set, and the (C)ECB all predict the same network structures, which are the efficient networks.[21]

The condition of strict top convexity is a strong condition. Also it is not necessary for each prediction of (C)ECB to be efficient under Y^{ce}, as the following example demonstrates.

Example 4. Let $N = \{1, 2, 3, 4\}$, $v(ij) = 4$ for $i \neq j$ and $v(g) = 0$ for any other network. In this example, v is not strictly top convex and each efficient network consists of two pairs. Furthermore, it is clear that under Y^{ce}, each (C)ECB predicts an efficient network.

5.4 *ECBs with a unique prediction*

Network formation games can be seen as a generalized version of characteristic function games. From Section 4.3, we know that static ECBs with a unique prediction completely characterize the core of a characteristic function game. One could wonder if a similar result can be obtained in network formation games.

It is not even immediate what core is in network formation games, because there might be widespread externalities across links. Nevertheless, we can show that it is possible to get the corresponding result with a pessimistic core definition.

Definition 20 (Pessimistic network core). Given a value function v and an allocation rule Y, the pessimistic network core is defined as

$$C(v, Y) = \{g \in \mathcal{G} \mid \nexists S \subseteq N \text{ and } g' \subseteq g^S \text{ such that } Y_i(g', v) > Y_i(g, v) \text{ for all } i \in S\}.$$

[21] Note that unlike the other results mentioned, Proposition 10 would not hold under top convexity. To see this, let $N = \{1, 2, 3\}$, $v(ij) = 4$ for $i \neq j$, $v(g^N) = 6$ for the complete network and $v(g) = 0$ for any other network. The value function is top convex. Nevertheless, it is easy to see that the networks with a single link can be supported as the prediction of a (C)ECB. The reason is that both the efficient network and the networks with a single link provide the same payoff to each individual getting a positive payoff.

The definition requires immunity from deviations to network structures that only include the deviators. This is the reason for the word "pessimistic": when deviating, a coalition does not just deviate to any other network, but contemplates only those networks that are isolated from the other players. Hence, further actions by other players cannot affect the deviators' payoffs.

Proposition 11 shows that if a certain network is predicted from every other network when the players are sufficiently patient, then this network is necessarily in the *pessimistic network core*. An ECB ϕ is said to have the unique absorbing network g^* if $\mathcal{T}(\sigma(g,h)) = g^*$ for every $(g,h) \in Z$, i.e., the path of play from every node terminates at g^*.

Proposition 11. *Suppose there exists δ^* such that for any $\delta \in (\delta^*, 1)$, ϕ is an ECB with the unique absorbing network g^*. Then $g^* \in C(v, Y)$.*

If g^* is the unique prediction from every node, then from any other network g', someone should be willing to take an action that eventually leads to g^*. For this, for any possible set of players S that form a component $c \in c(g')$, at least one person in S should (weakly) prefer g^* to g'; otherwise, the relevant component will never be disrupted. This shows that $g^* \in C(v, Y)$.

The converse is not true for one particular reason: the definition of the pessimistic network core ignores the dynamic nature of the problem considered here. That is, when checking for a deviation, it ignores checking for deviations in which the deviator cycles between several network structures. The next example demonstrates this.

Example 5. Let $N = \{1, 2, 3\}$ and let v be defined by $v(ij) = -2$, $v(ij, jk) = 9$, $v(ij, jk, ik) = -9$, and $v(\varnothing) = 0$. Finally let Y be anonymous and component balanced with $Y_j((ij, jk), v) = -3$. In this game, $C(v, Y)$ is composed of the empty network. However, we cannot support the empty network as the unique prediction of an ECB for δ large enough. The reason is that if the players deviate together and choose to cycle between the star structures, then each of them can get a higher payoff when δ is large enough.[22]

This problem does not exist in the static ECB, where the payoffs are realized only when a stable network forms. Indeed, it turns out that static ECBs with a unique prediction completely characterize the pessimistic network core.

With the static ECB, every ingredient of the extended coalitional game is the same, except for the preferences: Given a valuation function v and an allocation rule Y, the utility of player $i \in N$ from a finite path that terminates at some network g is $Y_i(g, v)$ and the player assigns the lowest utility to infinite paths.

Proposition 12.

- *If $g^* \in C(v, Y)$, then there exists a static ECB ϕ such that $\mathcal{T}(\sigma(g,h)) = g^*$ for every $(g,h) \in Z$.*
- *If ϕ is an ECB such that $\mathcal{T}(\sigma(g,h)) = g^*$ for every $(g,h) \in Z$, then $g^* \in C(v, Y)$.*

[22]The star network is a network where all players are linked to one central player and there are no other links. Here, it is $g = (ij, jk)$.

We have already seen the reason why if a network is the unique prediction, then it has to be in the pessimistic network core in the dynamic ECB when the discount factor approaches 1. The same reason also applies here. Hence, all we need to do is to show that each network in the pessimistic network core can be supported as the unique prediction of a static ECB.

Suppose that g^* is in the pessimistic network core. In any network $g \neq g^*$ in which there is a nonsingleton component $c \in g$ that is not a component of g^*, there necessarily exists an individual in that component who (weakly) prefers g^* to g; otherwise, g^* would not be in the pessimistic core. But then, at any network g that is not composed of the components of g^*, a player is willing to sever all of her links in expectation of g^*. The constructed coalitional behavior in the proof (see the Appendix) uses this fact and makes these players sever their links until we reach a network g' such that every nonsingleton component of g' is also a component of g^*. Following this, isolated players can merge to form g^*. Each profitable deviation from the described coalitional behavior is necessarily finite, as infinite paths provide the worst utility. Furthermore, it is easy to see that there is no one-step deviation, because at every node, the coalition that moves prefers the final outcome to the status quo and each path leads to the same final outcome. Hence, the described coalitional behavior is an ECB.

I complete this section with a few remarks on the modeling choices and how they affect the results.

Remark 2. I have mostly studied the dynamic ECB as opposed to the static ECB. This is without loss of generality. It is easy to see that all of the results still hold if we considered the static ECB with almost no change in the proofs.

Remark 3. I have chosen to model the network formation game as an infinite horizon game with a tree structure as opposed to a finite cyclic game. The choice of an infinite tree helps us capture the richness of the process by allowing us to consider coalitional behaviors in which actions taken at a network might depend on the coalitions that lead to that network. This would not be possible if we modeled network formation as a finite cyclic game and some results would change. For instance, the proof of Proposition 9 would no longer hold for the simple reason that it relies on individual specific punishments.

Remark 4. Finally, most of the literature considers pairwise deviations as opposed to the coalition-wise deviations that I consider. This decision is made for tractability. None of the results rely on the choice of coalition-wise deviations as opposed to pairwise deviations, but the choice of coalition-wise deviations greatly simplifies the argument and the proofs.

6. Literature review

In this section I briefly review the different strands of the literature on farsighted coalition formation. Inevitably, the review is incomplete; for extensive reviews, see Mariotti and Xue (2003), Ray (2008), and Ray and Vohra (2014).

The quest to incorporate foresight into the static cooperative solution concepts goes back at least to Harsanyi (1974), who criticized von Neumann and Morgenstern's stable set (1944) for being myopic. Chwe (1994) formalized Harsanyi's criticism and developed the solution concepts of the FSS and the LCS. These are set valued concepts in the tradition of von Neumann and Morgenstern's stable set, which use the indirect dominance relation instead of the direct dominance relation the stable set uses.[23]

Xue (1998) argued that this approach is not entirely satisfactory, as farsighted players should not only consider the final states their actions lead to, but they should also consider how these states are reached. By using Greenberg's (1990) framework, Xue (1998) proposed to use paths to incorporate foresight into his solution concepts.[24] But Xue (1998) still used a framework in which players form arbitrary expectations based on optimism or pessimism to evaluate different sets of paths. Questions remain about these extreme expectations players hold to evaluate sets of paths; for details, see Bhattacharya and Ziad (2012), Herings et al. (2004), and Ray and Vohra (2014).

Dutta and Vohra (2017) propose to deal with these issues by embodying the farsighted stable set with consistent expectations and introducing one-step deviations (also see Jordan 2006 for an earlier work concerning common expectations and farsighted stability, and see Bloch and van den Nouweland 2017, where players might hold heterogeneous expectations).[25]

Given the problems of the static approach, some authors found the way out by introducing an explicitly dynamic solution concept. The main solution concept in the dynamic approach is the EPCF (Konishi and Ray 2003 and Ray and Vohra 2014), which models coalition formation as an explicitly dynamic process, and the payoffs are discounted with a common discount factor.[26]

Other authors have tried to remedy the problems with the static approach by taking a cooperative domain such as the abstract game as the primitive and proposing a solution concept by defining a non-cooperative game from the abstract game. Examples include Herings et al. (2004) and Granot and Hanany (2016). Herings et al. (2004) take the abstract game as the primitive and define a multistage game associated with the abstract game; then they define an appropriate notion of rationalizability on this multistage game. In Granot and Hanany (2016), the evolution of play resulting from deviations is modeled as an extensive form game.

Recently, the literature has started to extend the solution concepts in two directions: (i) by incorporating history dependence and (ii) by allowing for different degrees of foresight. Vartiainen (2011) incorporates history dependent expectations into EPCF and shows that this greatly enhances the existence property of the solution concept. More

[23]For more on farsighted stable sets, see Diamantoudi and Xue (2003), Mauleon et al. (2011), and Ray and Vohra (2015). For more on the largest consistent set, see Béal et al. (2008), Bhattacharya (2005), Herings et al. (2009), Mauleon and Vannetelbosch (2004), Page et al. (2005), and Xue (1997).

[24]Also see Mariotti (1997) for an approach inspired by Greenberg (1990) on strategic form games.

[25]Another related work is Karos and Kasper (2018), in which the authors extend the work by Dutta and Vohra (2017) by using "extended expectation functions" to capture a coalition's belief about subsequent moves of other coalitions.

[26]Also see Dutta et al. (2005) and Vartiainen (2011).

recently, Dutta and Vartiainen (2020) and Ray and Vohra (2019) incorporate history dependent expectations to the REFS and SREFS. On the second development, Herings et al. (2019) introduced the concept of level-K farsightedness to incorporate differing levels of foresight. More recently, Herings et al. (2017) study matching when the players might have heterogeneous levels of foresight.[27]

7. A concluding remark

This section briefly describes the advantage of using CECB over LCRS with the favor exchange game analyzed in Section 4.1.

For convenience, I represent a coalitional behavior with three ordered letters that correspond to the actions chosen at each node. For instance, $\{F, N, F\}$ corresponds to the coalitional behavior where player 1 chooses F, and player 2 chooses N if he observes that player 1 has chosen F and chooses F if he observes N.

We already know that in this game, there is a unique CECB, which is $\{N, N, N\}$. This is not the case for the LCRS: it turns out that LCRS $= \{\{N, N, N\}, \{F, F, N\}, \{F, N, F\}\}$. To see the problem with the LCRS here, consider $\{F, F, N\}$. Player 2 has a profitable deviation from this coalitional behavior to $\{F, N, N\}$, but this deviation might be followed by a profitable deviation by player 1 to $\{N, N, N\}$, which would make player 2 worse off. Hence, player 2 will refrain from deviating in the first place and $\{F, F, N\} \in$ LCRS.

But there is a problem with this argument. In particular, player 2 does not choose the optimal action at one of his nodes, because he unreasonably believes that a deviation can trigger further deviations in the past that would make him worse off.

This problem means that LCRS will not posses the many desirable features of CECB such as backward induction, non-cooperative justification, and a relationship to certain attractive solution concepts such as subgame perfect equilibrium.

On top of this, we also see that LCRS has a multiplicity problem. Even in this simple example, we have multiple coalitional behaviors in the LCRS. CECB reduces these to a unique coalitional behavior. This suggests that CECB alleviates the problem of multiplicity without compromising from existence.[28]

Appendix

Proof of Lemma 1. Note that a credible set is a consistent set as in Chwe (1994) when one takes the set of coalitional behaviors as the set of outcomes and takes the profitable deviation relation as the effectiveness relation. Hence, the lemma directly follows by Proposition 1 in Chwe (1994). □

Proof of Proposition 1. *Step 1*: Showing the Existence for a Basic Game. Since each credible set can be rephrased as a consistent set (see the proof of Lemma 1), the existence in a basic game follows Proposition 2 in Chwe (1994).

[27] Also see Kirchsteiger et al. (2016), who provide evidence of behavior in favor of limited foresight.

[28] One might conjecture that if we refine LCRS by requiring that the solution concept should be an LCRS in every subgame, then we could get over these problems. Indeed, if we do this, then the problem in this example would disappear. Nevertheless, this solution concept would still suffer from similar drawbacks and we would still lose the certain desirable features of CECB such as backward induction.

Step 2: Generalizing to any Finite Game. We show that the CECBs of any finite game can be found through recursively finding the CECBs of basic games. The result will follow by Step 1.

Let Γ be any finite extended coalitional game. For any $z \in Z$, let $l(\Gamma(z))$ denote the length of the subgame at z, where length is the number of proper subgames $\Gamma(z)$ includes. Since the game is finite, there exists $z \in Z$ for which $l(\Gamma(z)) = 0$.

For any z for which $l(\Gamma(z)) = 0$, by Step 1, there exists a CECB for $\Gamma(z)$. Pick one CECB for each such z and call their union ϕ_0.

Suppose ϕ_i is defined for each $i < n$. For any z for which $l(\Gamma(z)) = n$, consider the reduced game at z given ϕ_{n-1}. Since this is a basic game, there exists a CECB for this game—call it ϕ'—and let $\phi_n = \phi_{n-1} \cup \phi'$. Continue this process until all the nodes are exhausted. The process defines an action for each node of the game and, hence, a coalitional behavior for the whole game.

Suppose ϕ is the resulting coalitional behavior. I show that ϕ is a CECB. Toward a contradiction, assume that there exists a profitable and credible deviation to ϕ'. Let $Z^0 = \{z \in Z \mid \phi(z) \neq \phi'(z)\}$ and let $Z^1 = \{z \in Z_0 \mid l(\Gamma(z)) \leq l(\Gamma(z'))\forall z' \in Z_0\}$. For any node $z \in Z^1$, the reduced game at z corresponds to the game used in the construction and, hence, by the construction there cannot exist a profitable and credible deviation from z in this reduced game. A contradiction. □

Proof of Proposition 2. Assume S has a profitable (and credible) deviation from ϕ to ϕ'. Let z^* be such that $\phi(z^*) \neq \phi'(z^*)$; $\sigma'(z^*)$ is not a cycle and there are finitely many $z \in \sigma'(z^*)$ such that $\phi(z) \neq \phi'(z)$. Then there exists $z' \in \sigma'(z^*)$ such that $\phi(z') \neq \phi'(z')$, but $\phi(z) = \phi'(z)$ for all $z \in \sigma'(z')$ such that $z \neq z'$. Think about the deviation in which S changes only $\phi(z')$ to $\phi'(z')$ and let the resulting coalitional behavior be ϕ''. This is a one-step deviation. Since the initial deviation is profitable, we have $\sigma''(z') \succ_S \sigma(z')$; hence, this one-step deviation is also profitable. Finally, credibility also follows as $\Gamma(z', \phi') = \Gamma(z', \phi'')$. □

The following definition and the lemma are useful in proving Propositions 3 and 4.

Definition 21. In an acyclic basic game with root z, we say that an action $a \in A_z$ s-dominates an action $(z, b, S) \in A_z$ if and only if there exists $i \in S$ with $a \succ_i (z, b, S)$.

An action $a \in A_z$ weakly s-dominates an action $(z, b, S) \in A_z$ if and only if there exists $i \in S$ with $a \succeq_i (z, b, S)$.

Lemma 2. *For an acyclic basic game, if a set of actions V satisfies internal stability with the s-dominance relation and external stability with the weak s-dominance relation, then any $a \in V$ can be supported as the stationary subgame perfect equilibrium of the bargaining game.*

Proof. Consider the following strategies: $i \in N$ offers the most preferred action in V; $i \in N$ accepts (z, x, S) if $(z, x, S) \in V$. If $(z, x, S) \notin V$, then the last player accepts if (z, x, S) is better than any stable action and rejects otherwise. The second to last player accepts if the last player accepts and if (z, x, S) is better than any stable outcome, and rejects otherwise, and so on. Note that by weak external stability, any action that is not in V will be rejected.

The strategy is stationary and it leads to an outcome in V. Furthermore, any action in V can be supported as the outcome with these strategies with an appropriate order of players. Hence, all we need to show is that the strategy is indeed subgame perfect.

First consider the acceptance–rejection stage. Suppose i accepts (z, x, S). First assume that $(z, x, S) \in V$. Note that by rejecting, i can only induce actions in V; however, by internal stability for any $(z, y, T) \in V$, $(z, x, S) \succeq_i (z, y, T)$. Now assume that $(z, x, S) \notin V$, but then i accepts only the action if it is better than every stable action and everybody else responding after him accepts the action (by rejecting, i can induce only actions in V).

Now suppose i rejects some action (z, x, S). But then this can only be because the implemented action is better than (z, x, S) or someone after i will reject, and the eventual implemented action is (weakly) worse than that implemented when i rejects. Finally, it is easy to see that at the proposal stage everyone is taking the optimal action. □

Proof of Proposition 3. By backward induction, it is sufficient to prove the result for a basic game. Let Γ be a strongly acyclic basic game with root z and let ϕ be an ECB for such a game.

Let $A_0 = \phi$. Let A_1 be the set of actions that are weakly s-dominated by ϕ. Let $A_2 = \phi'$ be an action that is undominated in $A \setminus (A_0 \cup A_1)$. Note that ϕ is not s-dominated by ϕ', as otherwise there would exist a profitable deviation from ϕ to ϕ'.[29] Let A_3 be the set of actions in $A \setminus (A_0 \cup A_1 \cup A_2)$ that are weakly s-dominated by ϕ'. Let $A_4 = \phi''$ be an action that is undominated in $A \setminus (A_0 \cup A_1 \cup A_2 \cup A_3)$. Continue until the set of all coalitional behaviors is exhausted (by strong acyclicity and finiteness, the set will eventually be exhausted).

Let $V = \bigcup_{k \text{even}} A_k$. The result follows by Lemma 2, since V satisfies internal stability with the s-dominance relation and external stability with the weak s-dominance relation. □

Proof of Proposition 4. All solution concepts can be found through backward induction; hence, it is sufficient to show the result for a basic game. Let Γ be a basic game with root z. Let S^* be the potent coalition and let ϕ^* be the corresponding coalitional behavior.

Let $\Phi^* = \{\phi \in \Phi \mid \phi \sim_S \phi^*\}$. It is clear that Φ^* is the set of ECBs of the game. Furthermore, since from any other coalitional behavior there exists a profitable deviation to Φ^*, it is also the set of CECBs. Finally, any $\phi \in \Phi^*$ satisfies external stability with the weak s-dominance and trivially satisfies internal stability with s-dominance;[30] hence, by Lemma 2, it can be supported as a stationary subgame perfect equilibrium.

[29] Let us say $\phi' = (z, x, T)$ and $\phi = (z, y, S)$. We already know that ϕ' is not weakly s-dominated by ϕ, which implies that $(z, x, T) \succ_T (z, y, S)$. But then if ϕ is s-dominated by ϕ', clearly there exists a deviation from ϕ to ϕ'

[30] The statement does not claim that Φ^* satisfies internal stability. Indeed Φ^* need not satisfy internal stability. The statement claims that a singleton from the set Φ^* trivially satisfies internal stability.

Now all we need to show is that no other coalitional behavior can be supported as a stationary subgame perfect equilibrium. Toward a contradiction, assume that in a basic game, $(z, x, T) \notin \Phi^*$ can be supported as a stationary subgame perfect equilibrium. Note that $S^* \cap T \neq \varnothing$, but no one in $S^* \cap T$ rejects (z, x, T) and offers ϕ^*, although they prefer ϕ^*. This could only be because there exists some $i \in S^*$ who rejects ϕ^* in anticipation of some ϕ'. But then $\phi' \succeq_i \phi^*$, implying $\phi' \sim_S \phi^*$. But then it is better for someone in $S^* \cap T$ to reject (z, x, T) and offer ϕ', a contradiction. □

Proof of Proposition 5. Both of the backward induction algorithms in Section 3 reduce down to the well known way to find subgame perfect equilibrium through backward induction in finite extensive form games.

To show that ECB is always a subgame perfect equilibrium, suppose ϕ is a coalitional behavior that is not a subgame perfect equilibrium. Then there exists an individual $i \in N$ who can deviate to a coalitional behavior ϕ' such that $\sigma'(z^*) \succ_i \sigma(z^*)$ for some $z^* \in Z$. Let $Z_1 = \{z \in \sigma'(z^*) \mid \phi(z) \neq \phi'(z)\}$ and let $Z_2 = \{z \in Z_1 \mid \sigma(z) \succeq_i \sigma'(z)\}$. If $Z_2 = \varnothing$, then consider the deviation from ϕ by i that includes only the actions in Z_1, which is a deviation that increases the payoff of i at every node at which an action changes. If $Z_2 \neq \varnothing$, then let z' be the node in Z_2 that is closest to z^*. Consider the deviation by i from ϕ that involves only the actions at the nodes in Z_1 that are between z^* and z' (including z^*, not including z'). The resulting deviation increases the payoff of i at every node at which an action changes; hence, it is profitable. That is, ϕ is not an ECB. □

Proof of Proposition 6. Take Γ that satisfies no indifference and monotonicity of actions.

First suppose that V is an SREFS and ϕ is the coalitional behavior that supports it. I show that ϕ is a static ECB. First note that static ECB satisfies the one-step deviation property on this domain. This is because any profitable deviation necessarily leads to finite paths. Hence, it suffices to check one-step deviations.

Suppose there is a profitable deviation at some unstable x in which some i blocks the move. But then $v_i(x) > v_i(\mathcal{T}(\sigma(x)))$, which is a contradiction to (ES). Suppose there is a profitable deviation at some unstable x where i blocks an action leading to z and T takes an action leading to y. By monotonicity of actions, there is also a profitable deviation in which i blocks the action leading to z and $T \cup \{i\}$ takes an action leading to y. Furthermore, since the initial deviation is profitable, we have $v_j(\mathcal{T}(\sigma(y))) > v_j(\mathcal{T}(\sigma(x)))$ for all $j \in \{T \cup i\}$, which is a contradiction to (M). Finally suppose there is a profitable deviation at some stable outcome x by coalition S to an outcome z. But then $v_j(\mathcal{T}(\sigma(z))) > v_j(\mathcal{T}(\sigma(x)))$ for all $j \in S$, a contradiction to (IS). We have exhausted all possible one-step deviations; hence, ϕ is a static ECB.

Now suppose that ϕ is a static ECB. We show that $S(\phi)$ is an SREFS supported by ϕ. Note that at any state x with $\phi(x) = (x, x, \varnothing)$, we have that there does not exist y and S such that $(x, y, S) \in A_x$ and $v_i(\mathcal{T}(\sigma(y))) > v_i(\mathcal{T}(\sigma(x)))$ for all $i \in S$, as otherwise S has a profitable deviation. Hence (IS) is satisfied. Now take any state x for which $\phi(x) = (x, y, S)$ for some y and S. First note that $v_j(\mathcal{T}(\sigma(x))) > v_j(x)$ for all $j \in S$, since otherwise, by no indifference, there exists $i \in S$ for which $v_i(x) > v_i(\mathcal{T}(\sigma(x)))$, in which case i has a profitable deviation at x. Hence, (ES) is satisfied. Finally, if (M) is violated, then trivially there exists a profitable deviation and ϕ is not an ECB, so (M) is also satisfied. □

Proof of Proposition 7. Figure 7 shows that an EPCF may not be a dynamic ECB, so here I show only that each dynamic ECB is an EPCF. Let ϕ be an ECB.

Take any $x \in Z$ with $\phi(x) = (x, y, S)$, where $x \neq y$. First observe that $\sigma(x) \succ_S (x, x, \varnothing)$, as otherwise, by no indifference, there exists $i \in S$ for which $(x, x, \varnothing) \succ_i \sigma(x)$, who would have a profitable deviation in which she blocks the taken action. This also implies that $\sigma(y) \succ_S \sigma(x)$. Toward a contradiction, assume that there exists z with $(x, z, S) \in A_x$ and $\sigma(z) \succ_S \sigma(y)$. But then there exists a profitable deviation by S to z, a contradiction. So ϕ satisfies the first condition.

Now assume that x is such that there exists y, S with $(x, y, S) \in A_x$ and $\sigma(y) \succ_S \sigma(x)$. Toward a contradiction, assume $\phi(x) = (x, x, \varnothing)$. Then $\sigma(y) \succ_S (x, x, \varnothing)$, but this implies that $\{(x, y, S), \sigma(y)\} \succ_S (x, x, \varnothing)$. Then there is a profitable deviation by S to y, a contradiction. So ϕ also satisfies the second condition and it is an EPCF. □

Proof of Proposition 8. The construction used in the proof is similar to that found in Diamantoudi and Xue (2003) and Konishi and Ray (2003).

Step 1. If $(x^*, P^*) \in C(N, V)$, then there exists an ECB ϕ such that $\mathcal{T}(\sigma((x, P))) = (x^*, P^*)$ for all $(x, P) \in Z$.

The proof is by construction. I construct an ECB with the desired property.

Take any $(x^*, P^*) = (x^*, \{S_1^*, S_2^*, \dots, S_K^*\}) \in C(N, V)$. For any coalition S, let $\overline{S}$ denote the partition of S composed of singletons. Let $\phi(x^*, P^*) = ((x^*, P^*), (x^*, P^*), \varnothing)$.

Let $(x_t, P_t) = (x_t, \{S_1^*, \dots, S_t^*, \overline{\bigcup_{j=t+1,\dots,K} S_j^*}\})$, where $x_t(S_j^*) = x^*(S_j^*)$ for all $j = 1, \dots, t$ and $x_t(\{i\}) = \max v(\{i\})$ for any $i \in \bigcup_{j=t+1,\dots,K} S_j^*$. For any $t = 0, 1, \dots, K-1$, let $\phi(x_t, P_t) = ((x_t, P_t), (x_{t+1}, P_{t+1}), S_{t+1}^*)$.

For any $(x, P) \neq (x_t, P_t)$ for some $t = 0, 1, \dots, K$, let i be the player with the smallest index for which $x_i^* \geq x_i$ and $i \in S \in P$, where $|S| \geq 2$. Note that since $(x^*, P^*) \in C(N, V)$, such a player exists. Let $\phi((x, P)) = ((x, P), (x', P'), \{i\})$, where $x'(\{i\}) = \max v(\{i\})$ and $\{i\} \in P'$.

This completes the specification of ϕ. Note that ϕ is a coalitional behavior, as it assigns a unique action for each $z \in Z$. Furthermore, $\mathcal{T}(\sigma(z)) = (x^*, P^*)$ for all $z \in Z$. Now we need to show that ϕ is an ECB. As any deviation leading to an infinite path would lead to the lowest utility, we can restrict attention to one-step deviations.

As any one-step deviation at (x^*, P^*) leads to a cycle, there exists no profitable one-step deviation at (x^*, P^*). Assume that there is a profitable one-step deviation at some $(x, P) \neq (x^*, P^*)$. But if the deviating coalition is taking another action at (x, P), then the deviation is not profitable since it will again end up at (x^*, P^*). Then the deviating coalition is inducing no-action at (x, P), but since the coalition moving at (x, P) weakly prefers (x^*, P^*) to (x, P), this cannot be a profitable deviation. Contradiction. Hence, ϕ is an ECB.

Step 2. Suppose ϕ is an ECB such that $\mathcal{T}(\sigma((x, P))) = (x^*, P^*)$ for all $(x, P) \in Z$. Then $(x^*, P^*) \in C(N, V)$.

Toward a contradiction, suppose ϕ is an ECB under foresight such that $\mathcal{T}(\sigma(x, P)) = (x^*, P^*)$ for every $(x, P) \in Z$ but (x^*, P^*) is not in the core. Then there exists (x, P) such that $S \in P$ and $x_S > x^*(S)$. But ϕ induces a finite path from (x, P) to (x^*, P^*) and at some point someone in S is active on this path. Let (x', P') be the first node on the path for which some $j \in S$ is active and let j deviate by refusing to take the action. The resulting deviation is profitable, as j is getting $x'_j = x_j$ instead of x^*_j. A contradiction. □

Proof of Proposition 9. Take any anonymous v. We prove the result using Y^{ce}, which is anonymous, is component balanced, and permits limited transfers. Let g^* be the efficient network we are trying to support. Let $g^* = \{c_1^*, c_2^*, \dots, c_K^*\}$, where the components are ordered in decreasing order of utility according to Y^{ce}.

Any network g can be written as $g = \{c_1^*, c_2^*, \dots, c_L^*, c_1, c_2, \dots . c_F\}$, where the highest paying L components of g^* is preserved and the remaining components are written in any order.[31] Let $I(c)$ denote the set of players who are a part of component c.

Now I describe the punishment network for player i given a network $g = \{c_1^*, c_2^*, \dots, c_L^*, c_1, c_2, \dots, c_F\}$, denoted by $g^i(g)$. This network is used to punish i whenever there is a deviation at some node z, where the path of play at z terminates at g. To get the punishment network, I slightly modify the algorithm Jackson (2003) uses to find a pairwise stable network under Y^{ce}.

If $i \notin \bigcup_{j=1,\dots,L} c_j^* = I^*$, then $g^i(g)$ is some network that satisfies (i) $g^i(g) = \{c_1^*, c_2^*, \dots, c_L^*, c'_1, c'_2, \dots, c'_F\}$, (ii) for any $k = 1, \dots, F$, c'_k maximizes $Y_j^{\text{ce}}(c, v)$ over j and c on the population $N \setminus I^* \setminus \bigcup_{t=1,\dots,k-1} I(c'_t)$, and (iii) $i \in c'_F$. Note that by the anonymity of v, we can always put i in the last component.

If $i \in c_l^*$, where $l \in \{1, \dots, L\}$, then $g^i(g)$ is some network that satisfies (a) $g^i(g) = \{c_1^*, c_2^*, \dots, c_{l-1}^*, c'_1, c'_2, \dots, c'_F\}$, (b) for any $k = 1, \dots, F$, c'_k maximizes $Y_j^{\text{ce}}(c, v)$ over j and c on the population $N \setminus (\bigcup_{t=1,\dots,l-1} I^*(c_t)) \setminus (\bigcup_{t=1,\dots,k-1} I(c'_t))$, and (c) $i \in c'_F$. Here again, we need the anonymity of v to put i in the final component.

Now I construct the (C)ECB, which consists of an initial phase and a punishment phase for each player.

Initial phase: At z_0, N moves to g^*. Thereafter no player takes an action. Hence, the game stays at g^* indefinitely.

Punishment phase of player i given g: Suppose the punishment phase of player i given g is triggered and suppose $g^i(g) = \{c_1^*, c_2^*, \dots, c_{l-1}^*, c'_1, c'_2, \dots, c'_F\}$. First the players in c_1^* move to form c_1^* if it has not formed yet. Then the players in c_2^* moves to form c_2^*. This continues up to the point where c_{l-1}^* has formed. Then players in c'_1 moves to form c'_1 and so on. Once $g^i(g)$ forms, no player takes an action.

Transitions between phases: The game starts at the initial phase. If at node z' a coalition S deviates from the phase that is supposed to be played and moves to node z, and if the path of play that is supposed to be played at z' terminates at $g = \{c_1^*, c_2^*, \dots, c_L^*, c_1, c_2, \dots, c_F\}$, then at node z, the punishment phase of player i given g is triggered, where i is a player getting the highest payoff in $I^*(g) \cap S$ if $I^*(g) \cap S \neq \varnothing$; otherwise, i is the player with the lowest index in S.

[31] If the highest paying component of g^*, which is c_1^*, is not preserved in g, then $g = \{c_1, c_2, \dots, c_F\}$, where the components are written in any order.

This completes the description of the coalitional behavior. Now we show that it is a (C)ECB. For this, it is sufficient to show that the described coalitional behavior is immune to profitable deviations. Take any node z. At z, suppose that the coalitional behavior leads to network $g = \{c_1^*, c_2^*, \ldots, c_L^*, c_1, c_2, \ldots, c_F\}$.

First, note that there does not exist a profitable deviation at node z by coalition S with $S \cap I^*(g) = \varnothing$. To see this, simply observe that given $c_1^*, c_2^*, \ldots, c_L^*$, g is the network that maximizes the payoff of each player in c_1, so nobody in c_1 can be a part of the deviation. But then, given $c_1^*, c_2^*, \ldots, c_L^*, c_1$, g is the network that maximizes the payoff of the players in c_2 and so on.

Second, note that there does not exist a profitable deviation at z by coalition S with $S \cap I(c_1^*) \neq \varnothing$. To see this, observe that if there is such a deviation, then $i \in c_1^*$ will be punished with a network $g^i(g) = \{c_1, c_2, \ldots, c_F\}$. By the above observation, the profitable deviation cannot include any node at this punishment phase; hence, $g^i(g) = \{c_1, c_2, \ldots, c_F\}$ is indeed the network the deviation would lead to. By the efficiency of g^*, $u_i(g^*) = u_i(g) \geq u_i(g^i(g))$. If $u_i(g^i(g)) < u_i(g^*)$, then we can find a large enough δ that would make the deviation unprofitable. If $u_i(g^i(g)) = u_i(g^*)$, then g^* is already the network that maximizes i's payoff; hence, the deviation is not profitable.

Finally, observe that if at node z there does not exist a deviation by a coalition S with $S \cap I(c_j^*) \neq \varnothing$ for any $j = 1, \ldots, t-1$, then at node z there does not exist a profitable deviation by any coalition S with $S \cap I(c_t^*) \neq \varnothing$. To see this, suppose that at node z there does not exist a deviation by a coalition S with $S \cap I(c_j^*) \neq \varnothing$ for any $j = 1, \ldots, t-1$. If there is a deviation by S with $S \cap I(c_t^*) \neq \varnothing$, then $i \in c_t^*$ is punished with a network $g^i(g) = \{c_1^*, \ldots, c_{t-1}^*, c_1, c_2, \ldots, c_F\}$. By the above observation, the profitable deviation cannot include any node at this punishment phase; hence, $g^i(g) = \{c_1^*, \ldots, c_{t-1}^*, c_1, c_2, \ldots, c_F\}$ is indeed the network the deviation would lead to. By the efficiency of g^*, $u_i(g^*) = u_i(g) \geq u_i(g^i(g))$. If $u_i(g^i(g)) < u_i(g^*)$, then we can find a large enough δ that would make the deviation unprofitable. If $u_i(g^i(g)) = u_i(g^*)$, then g^* is already the network that maximizes i's payoff; hence, the deviation is not profitable.

With the final observation, all possible deviations are exhausted; hence, the described coalitional behavior is an ECB, implying that it is also a CECB. □

Proof of Proposition 10. Suppose v is strictly top convex and let $Y = Y^{\text{ce}}$. We first show that g^* is efficient if and only if for every $i \in N$, we have $Y_i(g^*, v) > Y_i(g, v)$ for any network $g \neq g^*$. That is, each efficient network is also a network that (strictly) maximizes the payoff of each player.[32] If for every $i \in N$, we have $Y_i(g^*, v) > Y_i(g, v)$ for any network $g \neq g^*$, then g^* is trivially efficient. For the other direction, suppose that g^* is efficient, but $Y_i(g^*, v) \leq Y_i(g', v)$ for some $i \in N$ and some network g'. Let S be the coalition that is in the same component as i in g'. Note that by strict top convexity, g^* is a connected network in N, so $Y_i(g^*, v) = Y_j(g^*, v)$ for each $j \in N$. But then $Y_j(g'v) = Y_i(g', v) \geq Y_i(g^*, v) = Y_j(g^*, v)$ for any $j \in S$, a contradiction to strict top convexity.

Now, suppose that ϕ is a CECB, but there exists $z = (g, h)$ such that $\mathcal{T}(\sigma(z)) \notin E(v)$, where $E(v)$ is the set of efficient networks. At every such node z, let S be the minimal

[32] This statement is no longer true with top convexity.

coalition that can deviate from z to an efficient network g'. First note that the prediction from g' is necessarily an efficient network, as otherwise, in the reduced game at g', any player who moves can deviate to g', which would be both profitable and credible. It is credible, because there can be no profitable deviation from g', which (strictly) maximizes each player's payoff.

But then, at the reduced game at node z, think about the deviation where S moves to g' instead of whatever the CECB specifies. Since g' (strictly) maximizes every player's payoff, the deviation is profitable. Furthermore, for the same reason, the deviation cannot be followed by a further deviation; hence, it is credible, which is a contradiction to the initial supposition that ϕ is a CECB. Hence, the prediction of every CECB is efficient, implying that the prediction of every ECB is also efficient. □

Proof of Proposition 11. Suppose ϕ is an ECB such that $\mathcal{T}(\sigma(g,h)) = g^*$ for every $(g,h) \in Z$, but $g^* \notin C(v,Y)$. Then there exists a coalition S and a network $g' \subseteq g^S$ such that $Y_i(g',v) > Y_i(g,v)$ for all $i \in S$. Since the prediction from every node is g^*, there exists a node $z = (g,h)$, where $g' \in C(g)$ and some $i \in S$ is taking an action at z.[33] Consider the deviation, where i refuses to move at z. Then the payoff i receives will be higher for δ sufficiently large, as she is getting $Y_i(g',v)$ instead of eventually getting $Y_i(g,v)$. Hence, the deviation is profitable. □

Proof of Proposition 12. *Step 1.* Suppose ϕ is an ECB such that $\mathcal{T}(\sigma((g,h))) = g^*$ for all $(g,h) \in Z$. Then $g^* \in C(v,Y)$. The proof of Proposition 11 applies line by line, so I do not replicate it here.

Step 2. If $g^* \in C(v,Y)$, then there exists a static ECB ϕ such that $\mathcal{T}(\sigma((g,h))) = g^*$ for all $(g,h) \in Z$. Take any $g^* \in C(v,Y)$. Let $C(g^*) = \{c_1^*, c_2^*, \dots, c_k^*\}$ be the partition of g^* to its components and let $\{S_1^*, S_2^*, \dots, S_k^*\}$ be the partition of N, where S_i^* corresponds to the players in the component c_i^* for each $i = 1, 2, \dots, k$. For any coalition S, let $\overline{S}$ denote the empty network on S. Let $\phi(g^*, h) = ((g^*, h), (g^*, h), \varnothing)$, that is, every node associated with g^* is stable.

Let $g_t = \{c_1^*, \dots, c_t^*, \overline{\bigcup_{j=t+1,\dots,K} S_j^*}\}$. For any $t = 0, 1, \dots, K-1$, let $\phi(g_t, h) = ((g_t, h), (g_{t+1}, \cdot), S_{t+1}^*)$. For any (g,h), where $g \neq g^*$ and $g \neq g_t$ for some $t = 0, 1, \dots, K$, let i be the player with the smallest index for which $Y_i(g^*, c) \geq Y_i(g, v)$ and i is connected in g. Note that since $g^* \in C(v,Y)$, such a player necessarily exists. Let $\phi((g,h)) = ((g,h), (g', \cdot), \{i\})$, where $\{i\}$ is isolated in g', i.e., i severs all of her links.

This completes the specification of ϕ. Note that ϕ is a coalitional behavior, as it assigns a unique action for each $z \in Z$. Furthermore, $\mathcal{T}(\sigma(z)) = (g^*, \cdot)$ for all $z \in Z$. Now we need to show that ϕ is an ECB. As any deviation leading to an infinite path would lead to the lowest utility, we can restrict attention to one-step deviations.

Any one-step deviation at (g^*, h) leads to a path that again ends at the network g^*; hence, such a deviation would not be profitable. Assume that there is a profitable one-step deviation at some (g,h), where $g \neq g^*$. But if the deviating coalition is taking another action, then the deviation is not profitable since it will again end up at a node

[33] Note that in any network in which one of the components is g', someone from S needs to take an action for that component to get disrupted.

where the associated network is g^*. Then the deviating coalition is inducing no-action at (g, h), but since the coalition moving at (g, h) weakly prefers g^* to g, this cannot be a profitable deviation either—a contradiction. Hence, ϕ is an ECB. □

References

Béal, Sylvain, Jacques Durieu, and Philippe Solal (2008), "Farsighted coalitional stability in TU-games." *Mathematical Social Sciences*, 56, 303–313. [703]

Bhattacharya, Anindya (2005), "Stable and efficient networks with farsighted players: The largest consistent set." Report, University of York. [703]

Bhattacharya, Anindya and Abderrahmane Ziad (2012), "On credible coalitional deviations by prudent players." *Social Choice and Welfare*, 39, 537–552. [703]

Bloch, Francis and Anne van den Nouweland (2017), "Farsighted stability with heterogeneous expectations." Report. [703]

Chwe, Michael Suk-Young (1994), "Farsighted coalitional stability." *Journal of Economic theory*, 63, 299–325. [670, 676, 681, 682, 687, 703, 704]

Diamantoudi, Effrosyni and Licun Xue (2003), "Farsighted stability in hedonic games." *Social Choice and Welfare*, 21, 39–61. [695, 703, 708]

Dutta, Bhaskar, Sayantan Ghosal, and Debraj Ray (2005), "Farsighted network formation." *Journal of Economic Theory*, 122, 143–164. [673, 695, 696, 698, 703]

Dutta, Bhaskar and Hannu Vartiainen (2020), "Coalition formation and history dependence." *Theoretical Economics*, 15, 159–197. [704]

Dutta, Bhaskar and Rajiv Vohra (2017), "Rational expectations and farsighted stability." *Theoretical Economics*, 12, 1191–1227. [671, 687, 690, 703]

Feldman, Allan M. (1974), "Recontracting stability." *Econometrica*, 42, 35–44. [695]

Grandjean, Gilles, Ana Mauleon, and Vincent J. Vannetelbosch (2011), "Connections among farsighted agents." *Journal of Public Economic Theory*, 13, 935–955. [700]

Granot, Daniel and Eran Hanany (2016), "Subgame perfect farsighted stability." Report. [670, 692, 693, 703]

Green, Jerry R. (1974), "The stability of Edgeworth's recontracting process." *Econometrica*, 42, 21–34. [695]

Greenberg, Joseph (1990), *The Theory of Social Situations: An Alternative Game-Theoretic Approach*. Cambridge University Press, Cambridge. [703]

Harsanyi, John C. (1974), "An equilibrium-point interpretation of stable sets and a proposed alternative definition." *Management Science*, 20, 1472–1495. [676, 703]

Herings, P. Jean-Jacques, Ana Mauleon, and Vincent J. Vannetelbosch (2000), "Social rationalizability." Center discussion paper 81. [692, 693]

Herings, P. Jean-Jacques, Ana Mauleon, and Vincent J. Vannetelbosch (2004), "Rationalizability for social environments." *Games and Economic Behavior*, 49, 135–156. [670, 688, 692, 693, 703]

Herings, P. Jean-Jacques, Ana Mauleon, and Vincent J. Vannetelbosch (2009), "Farsightedly stable networks." *Games and Economic Behavior*, 67, 526–541. [673, 695, 696, 698, 699, 703]

Herings, P. Jean-Jacques, Ana Mauleon, and Vincent J. Vannetelbosch (2017), "Matching with myopic and farsighted players." Report. [704]

Herings, P. Jean-Jacques, Ana Mauleon, and Vincent J. Vannetelbosch (2019), "Stability of networks under horizon-K farsightedness." *Economic Theory*, 68, 177–201. [704]

Jackson, Matthew O. (2003), "The stability and efficiency of economic and social networks." In *Advances in Economic Design*, 319–361, Springer, Berlin. [709]

Jackson, Matthew O. (2010), *Social and Economic Networks*. Princeton University Press, Princeton. [695]

Jackson, Matthew O. and Anne van den Nouweland (2005), "Strongly stable networks." *Games and Economic Behavior*, 51, 420–444. [699, 700]

Jackson, Matthew O. and Asher Wolinsky (1996), "A strategic model of social and economic networks." *Journal of Economic Theory*, 71, 44–74. [695, 698]

Jordan, James S. (2006), "Pillage and property." *Journal of Economic Theory*, 131, 26–44. [703]

Karos, Dominik and Laura Kasper (2018), "Farsighted rationality." Report. [703]

Kirchsteiger, Georg, Marco Mantovani, Ana Mauleon, and Vincent J. Vannetelbosch (2016), "Limited farsightedness in network formation." *Journal of Economic Behavior & Organization*, 128, 97–120. [704]

Kóczy, László Á and Luc Lauwers (2004), "The coalition structure core is accessible." *Games and Economic Behavior*, 48, 86–93. [694, 695]

Konishi, Hideo and Debraj Ray (2003), "Coalition formation as a dynamic process." *Journal of Economic Theory*, 110, 1–41. [671, 687, 691, 694, 695, 703, 708]

Mariotti, Marco (1997), "A model of agreements in strategic form games." *Journal of Economic Theory*, 74, 196–217. [703]

Mariotti, Marco and Licun Xue (2003), "Farsightedness in coalition formation." In *Endogeneous Formation of Economic Coalitions* (Carlo Carraro, ed.), 128–155, Edward Elgar, Cheltenham Glos. [702]

Mauleon, Ana and Vincent J. Vannetelbosch (2004), "Farsightedness and cautiousness in coalition formation games with positive spillovers." *Theory and Decision*, 56, 291–324. [703]

Mauleon, Ana, Vincent J. Vannetelbosch, and Wouter Vergote (2011), "Von Neumann–Morgenstern farsightedly stable sets in two-sided matching." *Theoretical Economics*, 6, 499–521. [695, 703]

Page, Frank H., Jr., Myrna H. Wooders, and Samir Kamat (2005), "Networks and farsighted stability." *Journal of Economic Theory*, 120, 257–269. [703]

Ray, Debraj (1989), "Credible coalitions and the core." *International Journal of Game Theory*, 18, 185–187. [695]

Ray, Debraj (2008), *A Game-Theoretic Perspective on Coalition Formation*. Oxford University Press, London. [702]

Ray, Debraj and Rajiv Vohra (2014), "Coalition formation." In *Handbook of Game Theory With Economic Applications*, volume 4, 239–326, Elsevier, Amsterdam. [670, 671, 672, 687, 702, 703]

Ray, Debraj and Rajiv Vohra (2015), "The farsighted stable set." *Econometrica*, 83, 977–1011. [694, 695, 703]

Ray, Debraj and Rajiv Vohra (2019), "Maximality in the farsighted stable set." *Econometrica*, 87, 1763–1779. [704]

Rosenthal, Robert W. (1972), "Cooperative games in effectiveness form." *Journal of Economic Theory*, 5, 88–101. [670, 687]

Sengupta, Abhijit and Kunal Sengupta (1996), "A property of the core." *Games and Economic Behavior*, 12, 266–273. [695]

Vartiainen, Hannu (2011), "Dynamic coalitional equilibrium." *Journal of Economic Theory*, 146, 672–698. [703]

von Neumann, John and Oskar Morgenstern (1944), *Theory of Games and Economic Behavior*. Princeton University Press, Princeton. [703]

Xue, Licun (1997), "Nonemptiness of the largest consistent set." *Journal of Economic Theory*, 73, 453–459. [703]

Xue, Licun (1998), "Coalitional stability under perfect foresight." *Economic Theory*, 11, 603–627. [670, 687, 688, 703]

Co-editor Dilip Mookherjee handled this manuscript.

Manuscript received 28 October, 2017; final version accepted 12 April, 2019; available online 6 October, 2019.