Dynamic Coalitional Equilibrium*

Hannu Vartiainen†

March 6, 2010

Abstract

We study coalition formation processess of Konishi and Ray (2003). It is shown that an absorbing and deterministic process of coalition formation that also forms an equilibrium - satisfies a coalitional one-deviation property - does exist if one allows the process to be history dependent. All such dynamic equilibrium processes of coalition formation are characterized. Absorbing outcomes of dynamic equilibrium processes are also identified. It is shown that they always constitute a subset of the largest consistent set of Chwe (1994). A procedure that identifies a dynamic equilibrium process of coalition formation in finite time is constructed.

Keywords: one-deviation principle, coalition formation, history dependence

JEL: C71, C72, C78.

1 Introduction

As discussed by Ray (2007), any general model of farsighted behavior in a one-shot coalition game suffers from the so called *prediction problem*: the eventual payoff from a blocking cannot be predicted unless one knows the future blocking behavior of the coalitions.[1] But as further blockings should be evaluated according to the same criterion as the original one, there may not be a concrete final step from which to start the analysis. This is why static solution concepts, e.g. the core, are in trouble: they restrict coalitional behavior in way that is against its fundamental

*I am grateful for the three referees for excellent comments. I thank the seminar audiences in Universitat Autonoma Barcelona, Maastricht University and PCRC workshop for useful comments, and Ariel Rubinstein, Jean-Jacques Herings, Hannu Salonen, and Juuso Välimäki for beneficial comments.

†Turku School of Economics and Yrjö Jahnsson Foundation. E-mail: hannu.vartiainen@tse.fi.

[1] Aumann and Maschler (1974) is a case in point.

dynamic nature. Much of the modern literature of one-shot coalition formation, originated by Rosenthal (1972) and further developed e.g. by Greenberg (1990), Chwe (1994) and Xue (1998), capture aspects of this dynamic coalition formation process. A recurrent problem with these models is, however, that they either suffer from existence problems or tend to be too permissive by accepting outcomes that are intuitively not plausible (see Barberà and Gerber, 2007, for a formal argument on this).[2]

Konishi and Ray (2003) (henceforth KR) develop a modification of the standard, one-shot coalition formation framework that solves the prediction problem in an elegant way. In their model, payoffs of the agents do not materialize just once but they accumulate over time depending on the state of play at different points of time. States are affected by coalitional moves and future payoffs are discounted. A key benefit of this approach is that it allows consistent and explicit modeling of the dynamics of farsighted coalitional behavior: current coalitional move is justified by the prediction of the coalitional behavior that the move induces. KR establish a natural coalitional solution - *equilibrium process of coalition formation* - that captures these considerations. The solution can be (roughly) interpreted as a coalition version of the one-deviation principle. Since the process of coalition formation defines, for any infinite stream of states, an intertemporal (discounted) payoff, the model avoids the prediction problem. Importantly, KR show that a (possibly) randomized equilibrium process of coalition formation always exists, and that any absorbing equilibrium process only implements outcomes in the largest consistent set of Chwe (1994). However, KR also demonstrate that randomization or cycling may be needed for the existence of a Markovian equilibrium process of coalition formation where the coalitional move is conditioned only one the current state.

In an infinite horizon model, cyclic or random equilibrium processes of coalition formation have a clear meaning. However, in the classic one-shot framework they are more difficult to interpret. On the one hand, in a one-shot game, optimality of a coalitional move cannot dependent on the hypothesis that payoffs following the move are random - a property that may be needed for a process to be optimal in an infinite horizon model. On the other hand, if blockings are interpreted as coalitional negotiation prior to a binding agreement, as is often the case in the context of one-shot coalitional games, it is not clear which outcome could the coalitions "agree upon" if there is no state in which the play stays permanently.

But both of the above concerns are warranted only if there does not exist an equilibrium process of coalition formation that is detereministic and absorbing.

[2]Other recent models of farsighted coalition formation include Ray and Vohra (1997), Mariotti (1997), Barberà and Gerber (2003), Bloch and Gomez (2006), Diamantoudi and Xue (2006), Conley and Konishi (2002), Page and Wooders (2006), Page et al. (2005). Relatedly, Chatterjee et al. (1993), Bloch (1996), Ray and Vohra (1999), and Gomes and Jehiel (2005) model coalition formation non-cooperatively via a protocol.

Such equilibria do not contain randomization and, more importantly, converge to an outcome of which the coalitions can be interpreted to agree upon. The aim of this paper is to show that a version of detereministic and absorbing equilibrium process of coalition formation does exist if one drops the assumption that the processes are Markovian, *i.e.*, allows coalition formation processes to be *history dependent*.[3]

Formally, we apply a modified version of the solution concept of KR to the classic one-shot framework where coalitional negotiations take place prior to the binding agreement. We look at deterministic and finitely terminating processes of coalition formation that implement an outcome in finite time. They can, in the framework of KR, interpreted as deterministic and absorbing processes of coalition formation: processes that do not contain randomization and nor cycles in the long run. However, we relax the Markovian restriction imposed by KR. To highlight the possibility that coalition formation process may now depend on the history, we dub the solution as *dynamic* equilibrium process of coalition formation. The solution coincides with the standard one-deviation principle when restricted to games where each coalition consists of a single player.

It is important to note that our focus on deterministic and absorbing coalition formation processes is not a restriction on players' rationality but rather on the solution concept: it is more diffcult to find a solution with the desired properties than it is without them. In fact, our model can be interpreted as an undiscounted version of the KR framework, reflecting "full farsightedness" of the players (see below).

Our main result is that, due to flexibility coming with history dependency, a dynamic equilibrium process of coalition formation always exists. We show this by first characterizing the basic structure of all dynamic equilibrium processes of coalition formation. Like Xue (1998), the characterization is provided directly in terms of possible play paths. By varying possible histories, any dynamic equilibrium process of coalition formation induces a collection of play paths called a *consistent path structure*. This characterization is complete in a sense that each consistent path structure also gives rise to a dynamic equilibrium process of coalition formation that induces this collection of paths. Hence consistent path structures define a subset of outcomes that are implementable in equilibrium. Indeed, we characerize explicitly the largest set of outcomes that are implementable in any equilibrum. The existence of a consistent path structure is obtained by applying techniques from the social choice theory.[4] [5]

[3]KR finds the issue of history dependent coalition formation processes interesting and difficult but do not analyze it in detail.

[4]In particular those introduced by Fishburn (1977) and Dutta (1988).

[5]Equilibrium binding agreement -solution of Ray and Vohra (1997) coincides with dynamic equilibrium process of coalition formation when restricted to the class of games where coalitions

To see how the history dependency works, let us consider the "roommate problem" in KR (Example 10). There are three players $\{1, 2, 3\}$, three choices $\{x, y, z\}$, and payoffs are

	x	y	z
Player 1	1	0	a
Player 2	a	1	0
Player 3	0	a	1

where $a \in (0, 1)$. Possible coalitional moves are depicted in the commuting diagram in Figure 1.a, where $x \rightarrow_{\{2,3\}} y$ means that coalition $\{2, 3\}$ may change status quo outcome x to y, etc..

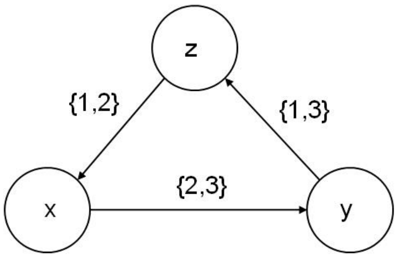

Figure 1.a.

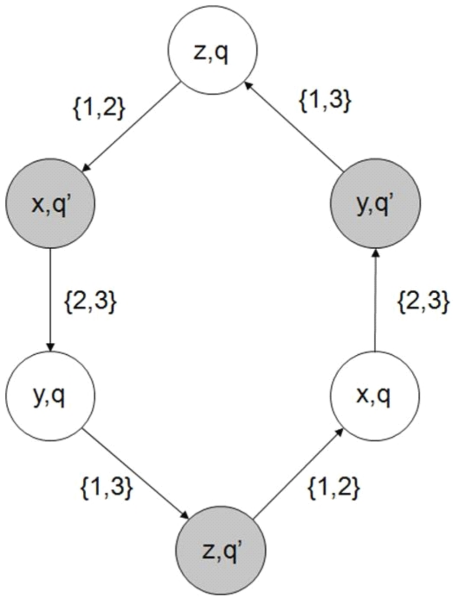

Figure 1.b.

The game proceeds as follows. At each point of time $t = 0, 1, ...$, there is a status quo outcome. This status quo can be replaced with a new outcome by the eligible coalition, as depicted in Figure 1.a. (e.g., $x \rightarrow_{\{2,3\}} y$ means x can be replaced with y by coalition $\{2, 3\}$). The new outcome becomes the status quo in $t + 1$. If the status quo outcome is not replaced with any other outcome, then it is implemented.

KR demonstrate that this game lacks an absorbing and deterministic equilibrium process of coalition formation, *i.e.,* a plan that specifies, at each node, whether the current status quo is implemented or not, and that is, at each node, optimal from viewpoint of the eligible coalition, given how the process behaves

can integrate but not disintegrate.

in the future (KR show, however, that the game does posses a randomized equilibrium process of coalition formation or, depending the value of paramter a, a cyclic equilibrium process of coalition formation). We argue that the game has an absorbing deterministic equilibrium process of coalition formation if one lets the process have memory. Let history affect the coalitional process via two *phases* $\{q, q'\}$, as follows. Divide the set of outcome-phase configuarions $\{x, y, z\} \times \{q, q'\}$ into two groups: terminal configuarions $\{(x, q), (y, q), (z, q)\}$ and transitory ones $\{(x, q'), (y, q'), (z, q')\}$. Let phase transitions from one configuration to another be determined as depicted in Figure 1.b: if the process moves away from a terminal configuration, then the phase changes from q to q', and if the process moves away from a transitory configuration, then the state changes from q' to q. Finally, let the phase dependent process of coalition formation implement the outcome on the table if the current configuration is terminal, and let the process move to the next configuration by changing the outcome on the table if the current configuration is transitory.

The constructed history dependent process of coalition formation is efficient in the sense that no coalition benefits from a one-time deviation to the process. In a terminal configuration, say (x, q), a devition leads to implementation of y via a transitory configuration (y, q'). This change is not profitable from the viewpoint of the deviating coalition $\{2, 3\}$. In a transitory configuration, say (y, q'), a deviation would lead to implementation of y (and not continuing to (z, q)), instead of z. Again, this change is not profitable from the viewpoint of the deviating coalition $\{1, 2\}$.

The underlying rationale for the existence of a history dependent equilibrium process of coalition formation can be seen from Figures 1.a,b. The commuting diagram in case (a) reflects a deterministic one-state Markov process, resulting in a cycle with even number of nodes. The second type (b) commuting diagram can in turn be interpreted as a two-state Markov process. A particular feature of a cycle with even number of nodes is that they are impossible to bipartition - divide into two disjoint sets such that every edge connects a node in the first set to one in the other. Bipartitioning is, however, possible when there are two phases. They double the number of outcome-phase nodes as in Fig. 1.b. And it is such bipartitioning on which one can build a (dynamic) equilibrium process of coalition formation: implement outcomes in the first set of nodes and transit from the other set to the first one.[6] [7]

In this example, any outcome could be implemented within the constructed coalition formation process depending on the initial outcome on the table. More

[6] Barberà and Gerber (2007) argue that there is no natural, well defined, internally consistent solution. Their argument assumes history *in*dependent coalitional strategies.

[7] That history dependence enlarges significantly possible bargaining outcomes is, of course, well known (see e.g. Chatterjee et al., 1993).

generally, however, the set of outcomes that are implementable in a dynamic equilibrium process of coalition formation constitute only a small subset of the outcomes - in fact, a subset of the largest consistent set of Chwe (1994).

There are two technical differences to the way the solution is formulated in KR (apart from history dependency). First, our focus is on properties of the solution that reflect coalitional optimality that is captured by the one-deviation principle. This is why we only require weak coalitional preference for transition from the current satus quo to the next one (as opposed to KR who require *strict* preference when the status quo outcome is not an efficient move for all coalitions). Second, we evaluate the intertemporal payoffs of the players by limit-of-the-means criterion rather than discounting. In the modeling of KR, this case can interpreted as the full farsightedness -benchmark. These differences notwithstanding, the key aspects of the solutions (*i.e.*, the one-deviation principle) are the same and warrant, in our view, common naming.

A problem with game theoretic solution concepts is often their computability. This is a particular concern here since the basis of our solution - the set of all finite play paths - is typically infinite. Hence it is not *a priori* clear whether there is any procedure that identifies an equilibrium in finite time. We show that this concern is not warranted. Our second key result is to construct an explicit procedure that identifies the outcomes that are implementable via any dynamic equilibrium process of coalition formation. Importantly, this procedure terminates in finite (bounded) time.

Relatedly, we also show that there is essentially no loss of generality in assuming that the processes of coalition formation are finitistic in a sense that (i) they have only finite memory (they are finite state Markov chains), and (ii) that they implement outcomes in bounded time. Thus starting from any history, there is an upper bound on the number rounds in which the outcome becomes implemented.

The Coase theorem famously asserts that, in absence of deadlines and other barriers on negotoation,interaction between agents always leads to efficiency. However, we demonstrate via an example that dynamic equilibrium processes of coalition formation are not always efficient (in our example every dynamic equilibrium process of coalition formation induces only inefficient outcomes).

There are interesting connections to other coalitional solutions concepts, in particular to the largest consistent set of Chwe (1994), conservative stable standard of behavior of Greenberg (1990) and Xue (1998), and full coalitional equilibrium of Mariotti (1997). Formal discussion of the connections to this literature is relegated to Section 4.

Section 2 defined the model and the solution concept. Section 3, which contains the main part of the paper, characterizes dynamic equilibrium processes of coalition formation, demonstrates its existence, and establishes a finite procedure to compute it. Section 4 compares the solution to the aforementioned solutions in the literature. Section 5 discusses the finitistic aspects of the solution. Section 6 comments briefly the efficiency of the solution, and Section 7 concludes with discussion.

2 Coalitional game

A *coalitional game* (see Rosenthal 1972, Greenberg 1990, Chwe 1994) is defined by a list

$$\Gamma = \langle N, X, (F_S)_{S\subseteq N}, (\succsim_i)_{i\in N}\rangle ,$$

where N is a finite set of players, X is a nonempty finite set of states or nodes or outcomes, an choice set $F_S(x) \subseteq X \cup \{\emptyset\}$ such that $\emptyset \in F_S(x)$ specifies the set of actions or effectiveness relations $F_S(x)$ available to a coalition $S \subseteq N$ at node $x \in X$.[8] Each player $i \in N$ has a preference relation $\succsim_i$ over the set of outcomes X.

The game is played in the following manner: There is an initial status quo x^0. At period $t = 0, 1, ...$, coalition S can challenge the current status quo outcome x^t by demanding an outcome $y \in F_S(x^t)$. In such case, y becomes the new status quo at period $t + 1$. If no coalition challenges x^t, then x becomes implemented. Only *one* coalition may be active at a time.

This model embodies several classic models, such as characteristic function games, games in strategic form or majority voting games. But it can also encompass games in partition function form (e.g. Ray and Vohra, 1997; Barberà and Gerber, 2003; Diamantoudi and Xue, 2006) or networks (e.g. Jackson and Wolinsky, 1996; Dutta and Mutuswami, 1997; Jackson and van den Nouweland, 2005; Page at al., 2006). For concrete examples of games falling into these categories, see Chwe (1994) or Ray (2007). In its general abstract form, the game has been analyzed by Chwe (1994), Xue (1997), and Konishi and Ray (2003). The origins of this modeling tradition were laid down by Greenberg (1990) and Rosenthal (1972).

In the remainder, we make the following assumption which simplifies the exposition: For all $x, y \in X$ such that $x \neq y$, there is at most one coalition S such that $y \in F_S(x)$. Such a coalition is denoted by $S(x, y)$. The only role of this assumption is to reduce the notational burden: when moving from x to y we do not need to specify which coalition induces the move.[9] To see why the assumption is without loss of generality, note that it actually only requires each outcome be indexed by the coalition that brought the outcome on the table (and the initial outcome

[8] Chwe (1994) does not assume finite outcome space.

[9] The assumption allows presenting the histories of play in terms of the nodes alone. Otherwise, a history should also specify which coalitions have been active along the play path.

is indexed by the empty set).[10] Hence the assumption makes no restrictions on the underlying physical structure. In particular, it does *not* affect any of the our results.[11]

Paths A *path* is a *finite* sequence $(x_0, ..., x_K)$ of outcomes such that $x_{k+1} \in \cup_S F_S(x_k)$, for all $k = 0, ..., K$. The *length* of the path $(x_0, ..., x_K)$ is K. Denote by X^k the $k-$fold Cartesian product of the set X, and denote the set of all paths by

$$\mathcal{X} = \cup_{k=1}^{\infty} X^k.$$

Further, denote the set of paths that start from node y by

$$\mathcal{X}_y = \{(x_0, ..., x_K) \in \mathcal{X} : x_0 = y\}.$$

For any collection $\mathcal{B} \subseteq \mathcal{X}$ of paths, denote the subcollection of paths that start from node y by

$$\mathcal{B}_y = \mathcal{B} \cap \mathcal{X}_y.$$

A path is abbreviated by $\bar{x} = (x_0, ..., x_K)$. By our expositional assumption, a path $\bar{x}$ also implicitly defines the coalitions that are active along the play. Denote the final element x_K of the path $(x_0, ..., x_K)$ by

$$\mu[(x_0, ..., x_K)] = x_K.$$

Process of coalition formation Denote the set of finite *histories* by $H := \mathcal{X}_{x^0}$. A deterministic *process of coalition formation* PCF σ is a function $\sigma : H \to X \cup \{\emptyset\}$. The interpretation of a PCF is that if $\sigma(h, x) = y \in X$, then the coalition $S(x, y)$ changes the status quo from x to y, and if $\sigma(h, x) = \emptyset$, then no coalition is active and x becomes implemented. Thus, a PCF specifies which - if any - coalition is active at the given history, and which outcome is the new status quo.

Denote, in the usual way, by $\sigma^t(h)$ the t^{th} iteration of σ starting from h, *i.e.*, $\sigma^0(h) = \sigma(h)$ and $\sigma^t(h) = \sigma(h, \sigma^0(h), ..., \sigma^{t-1}(h))$, for all $t = 1, ...$. A PCF σ is *terminating* if, for any $h \in H$ there is $T_h < \infty$ such that $\sigma^{T_h+1}(h) = \emptyset$. That is, after any history h, the process will implement the outcome $\sigma^{T_h}(h)$.

[10] If Y is the underlying physical outcome space, then we may let $X = Y \times 2^N$ such that $(y, S) \in X$ if and only if y is the physical outcome on the table and S is the most recent active coalition, and $S = \emptyset$ if $y = x^*$.

[11] The only consequence of the assumption is that some of the outcomes may have multiple representations in the set of outcomes. Since no restrictions are imposed on preferences over outcomes, this does not affect the characterization nor the existence results.

Let $\bar{\sigma}(h)$ denote the sequence of status quos in X that is induced by the strategy σ from the history h onwards

$$\bar{\sigma}(h) = (\sigma^0(h), \sigma^1(h), ...).$$

If σ is terminating, then $\bar{\sigma}(h)$ is finite and $\mu[\bar{\sigma}(h)]$ is well defined, for all h. Specifically, for a terminating PCF σ, if coalitional action $a \in X \cup \{\emptyset\}$ is chosen at history $(h, x) \in H$, then

$$\mu[\bar{\sigma}(h, x, a)] = \begin{cases} \mu[\bar{\sigma}(h, x, y)], & \text{if } a = y \in X, \\ x, & \text{if } a = \emptyset. \end{cases} \tag{1}$$

In particular, $\mu[\bar{\sigma}(h, \sigma(h))] = \mu[\bar{\sigma}(h)]$.

The solution Our primary question is whether equilibrium reasoning is compatible with the idea that a deterministic PCF is terminating. Our equilibrium condition, which is amended version of the solution in KR, is defined next.

Use the following notation for group preferences. For any $S \subseteq N$, and $x, y \in X$,

$$y \succ_S x \text{ if } y \succ_i x, \text{ for all } i \in S,$$
$$y \succsim_S x \text{ if } y \succsim_i x, \text{ for all } i \in S.$$

Take a PCF σ and a history (h, x). We say that coalition S has a *weakly preferred* move $y \in F_S(x)$ from x if $\mu[\bar{\sigma}(h, x, y)] \succsim_S x$. Further, a move $a \in F_S(x)$ is *efficient* for coalition S if there is no $b \in F_S(x)$ such that $\mu[\bar{\sigma}(h, x, b)] \succ_S \mu[\bar{\sigma}(h, x, a)]$.

Definition 1 (DEPCF) *A deterministic terminating PCF σ is a* dynamic equilibrium process of coalition formation *(DEPCF) if, for all $(h, x) \in H$,*

1. *$\sigma(h, x) \in X$ implies that $\sigma(h, x)$ is an efficient and weakly preferred move from x for coalition $S(x, \sigma(h, x))$.*
2. *$\sigma(h, x) = \emptyset$ implies that $\emptyset$ is an efficient move from x, for any coalition S.*

Thus a coalition is allowed to change the current status quo only if all its members agree on the move to the new status quo and cannot find any strictly better alternative status quo, given that the PCF is followed ever after. Moreover, if there is a strictly profitable move for some coalition, then the status quo must change.

Our solution reflects one-deviation principle: there is not profitable one-shot deviation for any coalition. The equilibrium condition in KR is slightly more stringent in that it requires that if there is a strictly profitable move for some coalition, then equilibrium move must also be *strictly* preferred for some (not necessarily the same) coalition.[12] (Our condition only requires the move to be weakly preferred). Discussion of the relation to KR and other solutions is relegated to Section 4.

Inducible paths and implementable outcomes In the remainder of this study, we characterize DEPCFs and verify their existence. Since the play path as well as the eventually implemented outcome may be sensitive to the choice of the initial status quo, which is usually somewhat arbitrary, it is natural to focus on *paths that are inducible* in DEPCF independently of how their first element has been reached. Our main interest is in outcomes that are *implementable* in DEPCF after some initial history.

More formally, paths that are induced in a PCF σ after some history of play are

$$\bar{\sigma}(H) = \{\bar{x} \in \mathcal{X} : \bar{\sigma}(h) = \bar{x} \text{ and } h \in H\}.$$

Denoting the set of final outcomes of a collection paths $\mathcal{B}$ by $\mu[\mathcal{B}] = \{\mu[\bar{x}] : \bar{x} \in \mathcal{B}\}$, the set of outcomes that are implementable in a terminating PCF σ is written as

$$\mu\left[\bar{\sigma}(H)\right] = \{x \in X : \mu[\bar{x}] = x \;\text{ and } \bar{x} \in \bar{\sigma}(H)\}.$$

Any outcome in $\mu\left[\bar{\sigma}(H)\right] \subseteq X$ - and nothing outside - can be implemented in a PCF σ by changing the initial history.

Simple recursive games In case there is only one potentially active coalition of a single player under each status quo, the game structure specifies a *simple recursive game* or, equivalently, a deterministic perfect information stochastic game with absorbing states. In this class of games, the player in a decision making turn faces the question of whether to stop the game or whether to move the decision making turn to one of the other players in his choice set. The essential problem is that of *commitment*: how to assign players a strategy that they all can commit to. This is not a trivial problem since no *a priori* restrictions are put on the players' preferences over possible outcomes.

When the game is equivalent to a simple recursive game, our equilibrium concept reduces to the standard form of the one-deviation principle: no player should benefit, after any history, from making a single deviation to his pure strategy.[13] Thus the dynamic equilibrium process of coalition formation can be seen as a coalitional generalization of the one-deviation property.

To our knowledge, there are no prior results concerning pure strategy subgame perfect equilibria in perfect in simple recursive games. Since a simple recursive

[12] That is, $\mu[\bar{\sigma}(h,x)] \succ_S x$, for some S.
[13] See e.g. Osborne and Rubinstein (1994).

game is a special case of a coalitional game á la Chwe (1994), our characterization and existence results below imply directly corresponding results in this class of stochastic games.

3 Characterization

The most primitive property of any equilibrium path $(x_0, ..., x_K) \in \mathcal{X}$ is *feasibility*:

$$x_K \succsim_{S(x_{k+1},x_k)} x_k, \text{ for all } k = 0, ..., K-1. \tag{2}$$

That is, a path is feasible if following the path is not worse for any member of active coalitions than stopping the game, provided that the final outcome of the path will be reached. Denote the set of feasible paths by $\mathcal{F} \subseteq \mathcal{X}$. Property 1 of Definition 1 hints that any coalitional equilibrium path is necessarily feasible. It will be convenient to work directly in terms of feasible paths.

Now we characterize coalitional equilibria in terms of the primitive data alone, *i.e.*, in terms of walkable paths. To this end, we define a dominance relation over paths. Recall that $\mathcal{F}_y$ is the set of feasible paths that originate from node y.

Definition 2 (Path Dominance) *A path* $\bar{y} \in \mathcal{F}$ dominates *a path* $\bar{x} \in \mathcal{X}$ *at the* k^{th} *step, denoted by* $\bar{y} \rhd_k \bar{x}$, *if* $y \in F_S(x_k)$, $\bar{y} \in \mathcal{F}_y$, *and* $\mu[\bar{y}] \succ_S \mu[\bar{x}]$, *for* $S = S(x_k, x_{k+1})$.

That is, if outcome x_k is reached along the path $\bar{x}$, then there is an active coalition at x_k that benefits from moving the play to path $\bar{y}$ under the hypothesis that the final element of $\bar{y}$ is implemented reached rather than that of $\bar{x}$.

Definition 3 (Consistent Path Structure) *A collection of feasible paths* $\mathcal{C} \subseteq \mathcal{F}$ *is a* consistent path structure *(CPS) if,*

(i) $\mathcal{C}_x$ *is nonempty, for all* $x \in X$,

(ii) for any $\bar{x} \in \mathcal{C}$, *there are* no k *and* y *such that* $\bar{y} \rhd_k \bar{x}$ *for all* $\bar{y} \in \mathcal{C}_y$.

A consistent path structure can be given a "blocking" interpretation as follows. (i) For any initial status quo node, there is a feasible path in the consistent path structure that guides how the play evolves. (ii) If an active coalition along the play path deviates from the path to a new node, then there is a path *in* the consistent path structure that starts from the new status quo and ends in an outcome that does *not* improve the payoffs of the members of the deviating coalition relative to what would be obtained if the original path is followed.

We now claim that a consistent path structure features stability in the sense described in the notion of dynamic equilibrium processes of coalition formation. The core of the argument is that any path of a consistent path structure is sustainable *if* any deviant coalition can be punished. But, by Definition 3 (ii), punishment *is* feasible if it is true that any path of a consistent path structure can be played. Hence an element of a consistent path structure is robust against one-shot (finitely many) deviations in a consistent way.

Intuitively, a consistent path structure does not aim to rule out outcomes with confidence. Rather, the idea is to characterize precisely the outcomes and paths that are consistent with equilibrium reasoning.

Now we verify the above argument formally, *i.e.,* that for any consistent path structure there exists a well defined PCF meeting the one-deviation property. Let $\mathcal{C}$ be a consistent path structure. Identify a function ξ on $\mathcal{C} \times \mathbb{N} \times X$ such that $\xi(\bar{x}, k, y) \in \mathcal{C}_y$ and $\xi(\bar{x}, k, y) \not\succ_k \bar{x}$, for all $(\bar{x}, k, y)$ such that $y \in F_{S(x_k, x_{k+1})}(x_k) \setminus \{x_k\}$. Since $\mathcal{C}$ satisfies Definition 3, such function does exist.

We now construct a deterministic and terminating PCF $\sigma^* : H \to X$ that is based on the function ξ. It is convenient to describe σ^* as a deterministic Markov chain $(\sigma^* : Q, g, \bar{x}^0)$, where Q is a set of *states* on which the process σ^* operates, g is a *transition function* from $Q \times X$ to Q, and $\bar{x}^0 \in \mathcal{C}_{x^0}$ is an *initial path* (which exists by Definition 3(i)). Later we shall argue that it without loss of generality to assume that the set of states Q is *finite*.

Let the set of states Q consist of pairs of paths and integers as follows:

$$Q = \{(\bar{x}, k) : \bar{x} = (x_0, ..., x_K) \in \mathcal{C} \text{ and } 0 \leq k \leq K\}. \tag{3}$$

Start with the path $\bar{x}^0$. Let the transition function g satisfy, for any $\bar{x} = (x_0, ..., x_K) \in \mathcal{C}$, for any $k = 0, ..., K-1$, and for any $y \in X$,

$$g((\bar{x}, k), y) = \begin{cases} (\bar{x}, k+1), & \text{if } y = x_{k+1}, \\ (\xi(x, k, y), 0), & \text{if } y \neq x_{k+1}. \end{cases} \tag{4}$$

Proceeding recursively from x^0, the set of histories H is partitioned by the set Q.

Let the PCF σ^* be conditional on the current state $(\bar{x}, k) \in Q$, where $\bar{x} = (x_0, ..., x_K)$, such that:

$$\sigma^*(\bar{x}, k) = \begin{cases} x_{k+1}, & \text{if } k < K, \\ \emptyset, & \text{if } k = K. \end{cases} \tag{5}$$

That is, the PCF continues along the path $\bar{x} = (x_0, ..., x_K)$ and implements x_K when the end of the path is reached.

Lemma 1 *A (deterministic and terminating) PCF σ^* is a DEPCF.*

Proof. Termination: Take $(q, y) \in Q \times X$ and let $g(q, y) = (\bar{x}, k) \in Q$, where $\bar{x} = (x_0, ..., x_K)$. Then, applying (4) and (5) recursively, σ^* implements an outcome in at most $K - k$ steps. Since (q, y) was an arbitrary element of $Q \times X$, σ^* is terminating.

Equilibrium condition (1): Take any path $\bar{x} = (x_0, ..., x_K)$ and state $(\bar{x}, k) \in Q$. Suppose that $\sigma^*(\bar{x}, k)$ is not an efficient move from x_k for coalition $S = S(x_{k+1}, x_k)$. Then there is $y \in F_S(x_k)\backslash\{x_{k+1}\}$ that induces a path $\xi(\bar{x}, k, y)$ such that $\mu[\xi(\bar{x}, k, y)] \succ_S \mu[\bar{x}]$. But by the definition of dominance, $\xi(\bar{x}, k, y) \rhd_k \bar{x}$ which contradicts part (ii) of a definition of CPS $\mathcal{C}$. Moreover, by construction, $\bar{x} \in \mathcal{F}$ which implies that also $\mu[\bar{x}] \succsim_S x_k$.

Equilibrium condition (2) is now obtained by replacing $k < K$ with $k = K$, coalition $S(x_{k+1}, x_k)$ with any coalition S, and $y \in F_S(x_k) \setminus \{x_{k+1}\}$ with $y \in F_S(x_K)$, and following the the first sentence of the proof of condition (1). ■

To fully characterize equilibrium strategies, consistent path structures need to be completed in the following sense: A consistent path structure $\mathcal{C}$ is *complete* if $(x_0, ..., x_K) \in \mathcal{C}$ implies $(x_k, ..., x_K) \in \mathcal{C}$, for all $k = 0, ..., K$. That is, following a path in a $\mathcal{C}$ is consistent with staying on a path in the $\mathcal{C}$. Note that completion is a purely expositional operation; existence of a complete $\mathcal{C}$ or its uniqueness is never an issue once $\mathcal{C}$ is specified.

Given a complete consistent path structure $\mathcal{C}$ and strategy σ^* defined on it, let $\bar{\sigma}^*(q)$ denote the path followed once $q \in Q$ has materialized. By construction,

$$\bar{\sigma}^*(\bar{x}, k) = (x_k, ..., x_K), \text{ for all } (\bar{x}, k) \in Q.$$

In particular, by the construction of σ^*,

$$\bar{\sigma}^*(\bar{x}, 0) \in \mathcal{C}.$$

Thus, by the definition of completeness of $\mathcal{C}$,

$$\bar{\sigma}^*(\bar{x}, k) \in \mathcal{C}, \text{ for all } (\bar{x}, k) \in Q.$$

Thus $\bar{\sigma}^*(Q) \subseteq \mathcal{C}$. Moreover, since $\bar{\sigma}^*(\bar{x}, 0) = \bar{x}$, for all $\bar{x} \in \mathcal{C}$, it follows that $\mathcal{C} \subseteq \bar{\sigma}^*(Q)$. We therefore have

$$\bar{\sigma}^*(Q) = \mathcal{C}. \tag{6}$$

That is, for any consistent path structure $\mathcal{C}$, we can find an equilibrium strategy that induces the corresponding complete $\mathcal{C}$

Now we prove the converse of Lemma 1 - that a consistent path structure characterizes behavior in any equilibrium, *i.e.*, that any collection of equilibrium paths is equivalent to a consistent path structure.

Lemma 2 *Let a (deterministic and terminating) PCF σ be a DEPCF. Then there is a complete CPS $\mathcal{C} \subseteq \mathcal{F}$ such that $\bar{\sigma}(H) = \mathcal{C}$.*

Proof. Take any $(h, x) \in H$. Then a finite path $\bar{\sigma}(h, x)$ exists since σ is terminating. We check both the defining conditions of CPS.

(i). By construction, $\bar{\sigma}(h, x) \in \mathcal{X}_x$. Since σ is a DEPCF, $\bar{\sigma}(h, x) \in \mathcal{F}$. Thus $\bar{\sigma}(h, x) \in \mathcal{F}_x$.

(ii). Let $\bar{\sigma}(h, x_0) = (x_0, ..., x_K) = \bar{x}$. First, take any $k < K$ and any $y \in F_S(x_k) \setminus \{x_{k+1}\}$, for $S = S(x_k, x_{k+1})$. Since x_{k+1} is efficient for S, it must be the case that $\mu[\bar{y}] \not\succ_S \mu[\bar{x}]$, where $\bar{\sigma}(h, x_0, ..., x_k, y) = \bar{y}$. That is, for any choice of k and y there is $\bar{y} \in \bar{\sigma}(H) \cap \mathcal{F}_y$ such that $\bar{y} \not\triangleright_k \bar{x}$.

Second, take $k = K$ and any $y \in F_S(x_K)$, for $S = S(x_K, y)$. Since $\emptyset$ is efficient for S, it must be the case that $\mu[\bar{y}] \not\succ_S \mu[\bar{x}] = x_K$, where $\bar{\sigma}(h, x_0, ..., x_k, y) = \bar{y}$. That is, for any choice of y there is $\bar{y} \in \bar{\sigma}(H) \cap \mathcal{F}_y$ such that $\bar{y} \not\triangleright_K \bar{x}$.

Finally, we argue that the CPS is also complete. Let $(x_0, ..., x_K) = \bar{\sigma}(h)$. It suffices to show that $(x_1, ..., x_K) = \bar{\sigma}(h, x_0)$. But this follows from the recursive structure of $\bar{\sigma}(h, x_0) = (\sigma(h), x_1, ..., x_K)$. ■

One the one hand by Lemma 1 and condition (6), any consistent path structure can be supported by a dynamic equilibrium process of coalition formation. On the other hand, by Lemma 2, a dynamic equilibrium process of coalition formation induces behavior consistent with a consistent path structure. We compound these observations in the following characterization.

Theorem 1 *A a (deterministic and terminating) PCF σ is an DEPCF if and only if there is a complete CPS $\mathcal{C}$ such that $\mathcal{C} = \bar{\sigma}(H)$.*

Theorem 1 permits us also to characterize the outcomes that are implementable within an equilibrium. The set of outcomes that are implementable within a consistent path structure $\mathcal{C}$ coincide with the set of outcomes that are inducible via an equilibrium and, conversely, the set of outcomes that are implementable within an equilibrium coincide with the set of outcomes that inducible within a $\mathcal{C}$. Summing these results gives a useful corollary.

Corollary 1 *There is a DEPCF σ such that $\mu[\bar{\sigma}(H)] = B$ if and only if there is a CPS $\mathcal{C}$ such that $\mu[\mathcal{C}] = B$.*

This result does not, however, say anything about the existence of a consistent path structure nor how it can be identified. The next section identifies the maximal consistent path structure, and shows that it always exists. This guarantees the existence of a solution.

3.1 Existence

The aim of this subsection is to prove that a consistent path structure and, hence, that a well defined coalitional equilibrium does exist. To this end, we need to define the following relation between paths and nodes. The concept is inspired by its cousin in the social choice literature (cf. Fishburn, 1977; Miller, 1980; and Dutta, 1988. Laslier, 1991, is an in-depth survey).

Let $\mathcal{B}$ be a set of paths.

Definition 4 *Path $\bar{x} \in \mathcal{X}$ is* covered *in $\mathcal{B}$ via node y if there is k such that $\bar{y} \rhd_k \bar{x}$, for all $\bar{y} \in \mathcal{B}_y$.*

That is, a path $\bar{x}$ is covered via node y in set $\mathcal{B}$ of paths if, at some particular node (the "k^{th}") of $\bar{x}$, the members of the active coalition profit by directing the play to node y rather than continuing along $\bar{x}$, no matter which paths in the set $\mathcal{B}$ set is followed after the deviation (under the hypothesis that the final element of played path is implemented).

Denote by $uc(\mathcal{B})$ the subset of elements in $\mathcal{B}$ that are *uncovered* in $\mathcal{B} \subseteq \mathcal{X}$, *i.e.*,

$$uc(\mathcal{B}) = \{\bar{x} \in \mathcal{B} : \bar{x} \text{ is not covered in } \mathcal{B}\}.$$

By construction, $uc(\mathcal{B}) \subseteq \mathcal{X} \cup \{\emptyset\}$.

We now strengthen the concept by iterating the uncovered-operator until no paths are left to be covered. Set $UC^0 = \mathcal{F}$, and let $UC^{t+1} = uc(UC^t) \subseteq \mathcal{F}$, for all $t = 0, \ldots$. The *ultimate uncovered set* $UUC \subseteq \mathcal{F}$ is then defined by $UUC := UC^\infty$.

Our aim is to show that UUC is a consistent path structure. That is, in addition to conditions stated in Definition 3, we need to verify that UUC is is a well defined concept. This automatcally the case if only finitely many itearations are needed.

Lemma 3 *There is $T < \infty$ such that $UC^T = UUC$.*

Proof. Call a family of paths

$$\left\{(x_0, ..., x_K) \in \mathcal{X} : \{(x_k, x_{k+1})\}_{k=0}^{K-1} = B \text{ and } x_K = y\right\}$$

a *dominance class,* parametrized by $B \subseteq X \times X$ and $y \in X$. Since X is a finite set, the cardinality of distinct dominance classes is finite, and they partition $\mathcal{X}$.

A dominance class contains all the relevant information concerning dominance: If two paths $\bar{x}$ and $\bar{x}'$ belong to the same dominance class, then $\bar{x}$ is covered in UC^t if and only if $\bar{x}'$ is covered in UC^t, for any t. This follows directly from the definitions of covering and dominance class. Since all paths in the same dominance class become covered at the same covering round t, and since there are finitely many dominance classes, the number of covering rounds to reach UUC must be finite. ■

Now we proceed to verify that UUC meets part (i) of Definition 3.

Lemma 4 *UC_x^t is nonempty, for all $x \in X$ and for all $t = 0, 1, ...$.*

Proof. Take any $S \subseteq N$, $x \in X$, and $t \in \{0, 1, ...\}$. Denote

$$C_S(x,t) = \left\{ y \in F_S(x) : (x) \text{ is covered in } UC^t \text{ via } y \right\}.$$

Further, let $D_S(x,t)$ contain any $y \in C_S(x,t)$ having the property that for any $\bar{y} \in UC_y^t$ and for any $z \in C_S(x,t)$ there is $\bar{z} \in UC_z^t$ that is not preferred to $\bar{y}$ for all members of S :

$$D_S(x,t) = \left\{ y \in C_S(x,t) : \bar{y} \in UC_y^t \text{ and } z \in C_S(x,t) \text{ imply there is } \bar{z} \in UC_z^t \text{ s.t. } \mu[\bar{z}] \not\succ_S \mu[\bar{y}] \right\}. \quad (7)$$

By the transitivity of $\succsim_S$, $D_S(x,t) = \emptyset$ if and only if $C_S(x,t) = \emptyset$.

Claim 0: Let $y \in D_S(x,t)$. Then $(x, \bar{y})$ is not covered at $k = 0$ in UC^t, for any $\bar{y} \in UC_y^t$.

Proof: If $(x, \bar{y}) \in UC_y^t$ is covered at $k = 0$ in UC^t, then there is $z \in F_{S(x,y)}(x)$ such that

$$\mu[\bar{z}] \succ_{S(x,y)} \mu[\bar{y}], \text{ for all } \bar{z} \in UC_z^t. \quad (8)$$

Since $y \in D_{S(x,y)}(x,t) \subseteq C_{S(x,y)}(x,t)$ we have $\mu[\bar{y}] \succ_{S(x,y)} x$. This means that (x) is covered in UC^t via z. But together with (8) this contradicts (7).□

Denote

$$G(x,t) = \cup_S D_S(x,t).$$

Claim 1: $G(x,t) = \emptyset$ if and only if $(x) \in UC^{t+1}$, for any $x \in X$ and for any $t = 0, 1, ...$.

Proof: By the definition of covering, if (x_J) is covered in UC^τ for any $\tau \leq t$, then it is covered in UC^t.

Call a path $(x_0, ..., x_J)$ a $G-$*path under* t if $x_{j+1} \in G(x_j, t)$ for all $j = 0, ..., K-1$. If, moreover, $G(x_J, t) = \emptyset$, then $(x_0, ..., x_J)$ is a *terminal* $G-$path under t.

Claim 2: For any $t = 0, 1, ...$, if $(x_0, ..., x_J)$ is a terminal $G-$path under t, then $(x_0, ..., x_J) \in UC^{t+1}$.

Proof: Let $(x_0, ..., x_J)$ be a terminal $G-$path under t. By Claim 0, $G(x_J, t) = \emptyset$ implies that $(x_J) \in UC^{t+1}$. Since $x_{j+1} \in G(x_j, t)$, for all $j = 0, ..., J-1$, it follows by using Claim 1 and iterating backwards on $j = 1, ..., J$ that $(x_0, ..., x_J) \in UC^{t+1}$.□

Claim 3: For any $t = 0, 1, ...$ and for any $x \in X$, there is a terminal $G-$path $(y_0, ..., y_J)$ under t such that $y_0 = x$.

Proof: By Claim 2, the claim implies that $(y_0, ..., y_J) \in UC^{t+1}$. Let, on the contrary of the claim, there be t that is the first stage in which the claim does not hold. For any x, denote $G^0(x,t) = \{x\}$, and define recursively

$$G^n(x,t) = \left\{ z : z \in G(y,t) \text{ and } y \in G^{n-1}(x,t) \right\}, \quad \text{for any } n = 1, ...$$

Since the claim does not hold in stage t, there is a particular x such that $G(z,t) \neq \emptyset$, for all $z \in G^n(x,t)$, for all n. Viewing G as the support of a Markov process with state space X, denote by $V \subseteq X$ the ergodic set of this process, i.e., V is the unique minimal subset of X in the sense of set inclusion such that $V = \cup_{v \in V} G(v,t)$ and $G^n(x,t) \subseteq V$ for all $n \geq n_V$ for some n_V. We now show a contradiction by proving that V cannot exist.

Subclaim 3.1: We first argue that $(v) \notin UC^t$ for all $v \in V$. Suppose on the contrary that $(v) \in UC^t$ and $v \in V$. By the definition of V, there is a $G-$path $(v_0, ..., v_L)$ under t such that $v = v_0 = v_L$. By using Claim 1 and iterating backwards on $\ell = 1, ..., L-1$, it follows that $(v_1, ..., v_L) \in UC^t$. But then, since $v_0 = v_J$, (v_0) is not covered in UC^t via v_1, a contradiction.

Subclaim 3.2: Now we argue that the opposite of Subclaim 3.1 cannot be true either. Suppose that $(v) \notin UC^t$ for all $v \in V$. Take any element of V, say y_0. Since $(y_0) \notin UC^t$, we have $G(y_0, t-1) \neq \emptyset$. By the maintained assumption, there is a terminal $G-$path $(y_0, ..., y_J)$ under $t-1$ such that, by Claim 2, $(y_0, ..., y_J) \in UC^t$. By Claim 1, $y_1 \in G(y_0, t-1)$ and $\bar{z} \in UC^t_{y_1} \subseteq UC^{t-1}_{y_1}$ imply $(y_0, \bar{z}) \in UC^t_{y_0}$. By the definition of V, there is a $G-$path $(v_0, ..., v_L)$ such that $y_0 = v_0 = v_L$. Thus by using Claim 1 and iterating backwards on $\ell = 0, ..., L-1$, it follows that $(v_0, ..., v_{L-1}, \bar{z}) \in UC^{t+1}_{y_0} \subseteq UC^t_{y_0}$ for all $\bar{z} \in UC^t_{y_1}$. This implies, since $v_1 \in G(y_0, t)$, that also $y_1 \in G(y_0, t)$. By the definition of V, then, $y_1 \in V$. Itarating this way on $j = 2, ..., J-1$ it follows that $y_J \in V$. But by $(y_0, ..., y_J) \in UC^t$ it also follows that $(y_J) \in UC^t$, and a contradiction is proved.□

Claims 2 and 3 now establish the proof. ■

Finally, we state our existence result: the ultimate uncovered set UUC is a consistent path structure.[14] For this, we need to show that UC^T_y is nonempty for all $y \in X$ and that no element in UUC is covered in UUC.

Theorem 2 *UUC is a CPS.*

Proof. By construction, $UUC \subseteq \mathcal{F}$. Thus it suffices to check parts (i) and (ii) of Definition 3.

[14] Recently, and independently of this study, Flesch et al. (2008) achieve a closely related existence result concerning subgame perfect Nash equilibria in perfect information stochastic games.

(i). By Lemma 3, $UC_y^T = UUC_y$ which is, by Lemma 4, nonempty, for all $y \in X$.

(ii). By the construction of UUC, $uc(UUC) = UUC$. Thus $\bar{x} \in UUC$ is not covered in UUC, *i.e.*, there is no y and k such that $\bar{y} \rhd_k \bar{x}$, for all $\bar{y} \in UUC_y$. ■

The next result shows that UUC is the (unique) maximal consistent path structure in the sense of set inclusion.

Theorem 3 *UUC contains as a subset any CPS.*

Proof. Let $\mathcal{C}$ be a CPS. Take any $\bar{x} \in \mathcal{C}$.

1^{st} iteration: Since $\bar{x}$ satisfies Definition 3(ii), and $\mathcal{C} \subseteq \mathcal{F}$, it follows that $\bar{x}$ is not covered in $UC^0 = \mathcal{F}$. Hence $\bar{x} \in uc(\mathcal{F}) = UC^1$. Since $\bar{x}$ was arbitrary, $\mathcal{C} \subseteq UC^1$.

2^{nd} iteration: Since $\bar{x}$ satisfies Definition 3(ii), and $\mathcal{C} \subseteq UC^1$, it follows that $\bar{x}$ is not covered in UC^1. Hence $\bar{x} \in uc(UC^1) = UC^2$. Since $\bar{x}$ was arbitrary, $\mathcal{C} \subseteq UC^2$.

$\vdots$

T^{th} iteration: Since $\bar{x}$ satisfies Definition 3(ii), and $\mathcal{C} \subseteq UC^{T-1}$, it follows that $\bar{x}$ is not covered in UC^{T-1}. Hence $\bar{x} \in uc(UC^{T-1}) = UC^T$. Since $\bar{x}$ was arbitrary, $\mathcal{C} \subseteq UC^T =: UUC$. ■

Since the UUC is obtained via a well defined recursive process, there is no question about its existence. By Theorems 2 and 1, one can construct a terminating equilibrium process of coalition formation on UUC. Hence we can conclude that a dynamic equilibrium process of coalition formation is guaranteed to exist.

Corollary 2 *There is a DEPCF σ such that $\bar{\sigma}(H) = UUC$.*

By Theorem 3, we have also characterized outcomes that are implementable with any equilibrium.

Corollary 3 *The set of outcomes that are implementable via any DEPCF is contained in $\mu[UUC]$.*

3.2 Computational Considerations

The problem with game theoretic solution concepts is often their computability. Such questions are particularly acute here since the set of paths $\mathcal{X}$ typically contains infinitely many elements. It should not be clear *a priori* whether coalitional equilibria, or outcomes they implement, can be identified via a procedure that terminates in finite time. Applicability of such a solution would, of course, be questionable.[15]

We now show that these concerns are unwarranted. A convenient algorithm for computing the relevant elements of the ultimate uncovered set, *i.e.* the largest consistent path structure, is provided. This algorithm, which terminates in finite time, thus generates a complete description of the outcomes that can be implemented via a dynamic process of coalition formation.

In order to establish the desired results, we develop a notion that captures the relavent information contained in a path in the simpliest possible form. We say that $\bar{y} = (y_0, ..., y_L) \in \mathcal{X}$ is a *reduction* of $\bar{x} = (x_0, ..., x_K) \in \mathcal{X}$ if $x_0 = y_0$, $x_K = y_L$, and $\{(y_l, y_{l+1})_{l=0}^{L-1}\} \subseteq \{(x_k, x_{k+1})_{k=0}^{K-1}\}$. Then $\bar{y}$ is a *full reduction* of $\bar{x}$ if it is a reduction of $\bar{x}$ and if the only reduction of $\bar{y}$ is $\bar{y}$ itself. That is, $\bar{y}$ contains a minimum amount of edges of $\bar{x}$ that are needed to travel from x_0 to x_K. Note that the reduction -relation is transitive. Since the paths contain at most finitely many distinct elements, each $\bar{x}$ has a full reduction. However, a full reduction of $\bar{x}$ need not be unique.

For any set $\mathcal{B}$ of paths, denote by $fr(\mathcal{B})$ the collection of *all* full reductions of the elements in $\mathcal{B}$, *i.e.*, the elements of $\mathcal{X}$ that are full reductions of $\mathcal{B}$. Note that any path $(x_0, ..., x_K)$ that contains a cycle such that $x_k = x_l$ for some $k < l$, cannot be a full reduction since $(x_0, ..., x_k, x_{l+1}, ..., x_K)$ is a reduction of $(x_0, ..., x_K)$ but not vice versa. Thus, since any fully reduced path is acyclic and since X contains finitely many elements, the set of fully reduced paths $fr(\mathcal{X})$ can be identified in finite time.

The full reduction -operation preserves two important aspects of consistent paths structures: the initial status quo and the feasibility. More formally,

$$\mathcal{B} \subseteq \mathcal{F}_x \text{ implies } fr(\mathcal{B}) \subseteq \mathcal{F}_x. \tag{9}$$

The following observation is now easily deduced from (9).

Lemma 5 *If $\mathcal{C}$ is a CPS meeting Definition 3 then so is $fr(\mathcal{C})$.*

Thus, when identifying consistent path structures - or outcomes that can be implemented via them - it is in general sufficient to focus on consistent path structures that are composed of fully reduced feasible paths $fr(\mathcal{F})$. In particular, the fully reduced form of the feasible ultimate uncovered set $fr(UUC)$ is a consistent path structure.

Now we describe a finite procedure that identifies the fully reduced ultimate uncovered set $fr(UUC)$. Identify $fr(\mathcal{F})$. Define recursively the uncovered set and its iterations on the set of fully reduced paths: $UC_{FR}^0 = fr(\mathcal{F})$ and $UC_{FR}^{j+1} =$

[15] See Dutta (1988) for a related and inspiring discussion.

$uc(UC_{FR}^{j})$ for all $j = 0, \ldots$. The fully reduced ultimate uncovered set is then defined by $UUC_{FR} = UC_{FR}^{\infty}$. Since, by the argument made in Lemma 3, only finitely many iterations are needed for UC_{FR}^{t} to converge, UUC_{FR} can be identified in finite time.

Thus UUC_{FR} can be identified in finite time. Our aim is to show that UUC_{FR} is a CPS and contains as subsets all fully reduced CPSs. This implies that the endpoints of the paths in UUC_{FR} are the outcomes that are implementable in any dynamic equilibrium process of coalition formation.

Our first task is to show that UUC_{FR} coincides with $fr(UUC)$ and is contained in UUC.

Lemma 6 $fr(UUC) = UUC_{FR} \subseteq UUC$.

Proof. *Claim 1*: $UC_{FR}^{t} \subseteq UC^{t}$, for all t.

Proof: Let t be the first stage in which $\bar{x} \in UC_{FR}^{t} \setminus UC^{t}$, for some $\bar{x}$. Hence $UC_{FR}^{t-1} \subseteq UC^{t-1}$. But then, by the definition of covering, since $\bar{x}$ is covered in UC^{t-1}, it must be covered in UC_{FR}^{t-1}. But this contradicts $\bar{x} \in UC_{FR}^{t}$.□

Claim 2: $fr(UC^{t}) \subseteq UC^{t}$, for all t.

Proof: Let t be the first stage in which there is $\bar{x} \in fr(UC^{t}) \setminus UC^{t}$. Since $\bar{x} \in fr(UC^{t})$, $\bar{x}$ must be a full reduction of some $\bar{y} \in UC^{t}$. But by the definition of full reduction, the assumption that $\bar{x}$ is covered in UC^{t-1}contradicts the assumption that $\bar{y}$ is not covered in UC^{t-1}.□

Claim 3: $UC_{FR}^{t} = fr(UC^{t})$ for all t.

Proof: Since the full reduction of a set of fully reduced paths is the set itself, it follows by Claim 1 that $UC_{FR}^{t} = fr(UC_{FR}^{t}) \subseteq fr(UC^{t})$, for all t. For the other direction, let t be the first stage in which $\bar{x} \in fr(UC^{t}) \setminus UC_{FR}^{t}$, for some $\bar{x}$. Hence, $fr(UC^{t-1}) \subseteq UC_{FR}^{t-1}$. But then, by the definition of covering, since $\bar{x}$ is covered in UC_{FR}^{t-1} it must be covered in $fr(UC^{t-1})$. By the definition of full reduction, $\bar{x}$ is also covered in UC^{t-1}. Hence $\bar{x} \notin UC^{t}$. But then $\bar{x} \in fr(UC^{t}) \setminus UC^{t}$, which contradicts Claim 2. Thus $fr(UC^{t}) \subseteq UC_{FR}^{t}$, for all t.□

Claim 4: $fr(UUC) = UUC_{FR} \subseteq UUC$.

Proof: Combining Claims 1 and 3, we have $fr(UC^{t}) = UC_{FR}^{t} \subseteq UC^{t}$ for all t. The fact that only finitely many iterations are needed gives the result. ■

From Lemmata 5 and 6 it now follows that also the largest consistent path structure can be described directly in terms of fully reduced paths via the iterative procedure that identifies UUC_{FR}.

Theorem 4 *UUC_{FR} is a CPS. Moreover, UUC_{FR} contains as a subset any fully reduced CPS.*

Proof. By Lemma 5, $fr(UUC)$ is a CPS. By Lemma 6, $fr(UUC) = UUC_{FR}$. Thus UUC_{FR} is a CPS.

By Theorem 2, UUC contains any CPS $\mathcal{C}$ as a subset. Thus, by Lemma 6, UUC_{FR} contains $fr(\mathcal{C})$ as a subset. ■

Furthermore, since the final element of a path is invariant with respect to the full reduction -operation, *i.e.*, $\mu[\mathcal{B}] = \mu[fr(\mathcal{B})]$ for any $\mathcal{B} \subseteq \mathcal{X}$, it follows that the outcomes that can be implemented with UUC coincide with the outcomes that can be implemented with UUC_{FR}.

Theorem 5 *The set of outcomes that are implementable via any DEPCF is contained in* $\mu[UUC_{FR}]$.

Proof. By Corollary 3, Lemma 6, and since $\mu[UUC_{FR}] = \mu[UUC]$. ■

We have thus constructed a procedure that specifies, in finitely many steps, the outcomes that can be implemented in *any* equilibrium process of coalition formation. The collection of paths identified by the procedure also contains the reduced forms of all dynamic equilibrium processes of coalition formation and, moreover, it constitutes itself a dynamic equilibrium process of coalition formation. It is fair to say that the procedure provides all the information that is needed for the analysis of coalitional behavior.

4 Finitistic Equilibria

We have assumed that the processes of coalition formation are history dependent and implement an outcome in finite time. These assumptions problematic if (a) the process has to have infinite memory, (b) there is no upper bound on how long will it take to implement an outcome. In this section we collect the results of the previous sections to argue that neither of these concerns is warranted.

Recall that the dynamic process of coalition formation σ^* used in Lemma 1 uses a state space Q that has the same cardinality as the consistent path structure $\mathcal{C}$ on which it is built. By Lemma 5, a full reduction $fr(\mathcal{C})$ of $\mathcal{C}$ is also a consistent path structure. In terms of real consequences, however, $fr(\mathcal{C})$ and $\mathcal{C}$ are equivalent, *i.e.,* $\mu[fr(\mathcal{C})] = \mu[\mathcal{C}]$. Since fully reduced paths are necessarily acyclic, it follows that $fr(\mathcal{C})$ contains finitely many elements. Since, by Theorem 1, any consistent paths structure can be identified with a dynamic equilibrium processes of coalition formation, we can conclude that, in terms real consequences (implemented outcomes) it is without loss of generality to assume that the process is finite state Markov chains.

Corollary 4 *There is a DEPCF that is a finite state Markov chain. Moreover, for any DEPCF there is another DEPCF, equivalent in terms of real consequences, that is a finite state Markov chain.*

Moreover, ince dynamic equilibrium processes of coalition formation can be identified with consistent path structures, any property of paths in a consistent path structure translated directly to an observation concerning equlibria. Since $fr(\mathcal{C})$ contains finitely many paths that are, at most, finitely long, final element of any path in $fr(\mathcal{C})$ can be reached in finitely many steps. Hence Theorem 1 implies the following corollary.

Corollary 5 *There is a DEPCF that that implements an outcome in uniformly bounded time after any history. Moreover, for any DEPCF there is another DEPCF, equivalent in terms of real consequences, that implements an outcome in uniformly bounded time after any history.*

5 Relation to Other Models

In this section, we relate our solution to some existing equilibrium notions of coalition formation.

5.1 Blocking in Real Time

In this section we relate our one-shot coalitional game to the model of real time blocking of KR (Konishi and Ray, 2003) (see also Ray, 2007).[16] Now X is is interpreted as a set of *states* and u_i is a utility function of player $i \in N$ over X. Analogously to the model above, $F_S(x) \subseteq X$ is the set of states achievable by a one-step coalitional move in state x by a coalition S. It assumed that $x \in \cap_S F_S(x)$ to guarantee that staying in x - once reached - is possible.

Let H be the set of all histories of states $(x_0, ..., x_t)$ such that $x_0 = x^*$. A deterministic process of coalition formation PCF is now a function $p : H \to X$, capturing transitions from one state to another. Letting $p^0(h) = p(h)$ and $p^t(h) = p(h, p^0(h), ..., p^{t-1}(h))$, for all $t = 1, ...$, denote by $\bar{p}(h) = (p^0(h), p^1(h), ...)$ the chain of states that will materialize along the PCF p, starting from history h. All walkable chains are denoted by $\bar{p}(H) = \{\bar{p}(h) : h \in H\}$. A PCF p is *absorbing*:[17]

[16] Hyndman and Ray (2007) and Gomez and Jehiel (2005) are other papers with real time coalitional negotiation. However, there the negotiation is noncooperative, based on a protocol.

[17] Equivalently, we assume that p is an absorbing (deterministic) Markov chain with finite state space.

if for any history h there is an integer $T_h < \infty$ such that there is $x \in X$ such that

$$p^t(h) = x, \text{ for all } t > T. \tag{10}$$

The absorbing state of the PCF p starting from history h, denoted by $\alpha[\bar{p}(h)]$, is then well defined for all h.

We evaluate the players' intertemporal payoffs by the *limit-of-the-means* criterion:

$$V_i(h) = \lim_{T\to\infty} \frac{\sum_{t=0}^{T} u_i(p^t(h))}{T}.$$

Such payoff exists whenever PCF p is absorbing:

$$V_i(h) = u_i(\alpha[\bar{p}(h)]), \quad \text{for all } h \in H.$$

Profitable moves and efficiency concepts are defined analogously to the one-shot game discussed before. That is, given an absorbing PCF p and a history (h, x), coalition S has a weakly preferred move $y \in F_S(x)$ from x if $V_i(h, x, y) \geq u_i(x)$ for all $i \in S$. Further, a move $y \in F_S(x)$ is efficient for coalition S if there is no $z \in F_S(x)$ such that $V_i(h, x, z) > V_i(h, x, y)$ for all $i \in S$.

Given these notions, we may now define a modifed version of the equilibrium criterion of KR. A dynamic, deterministic, absorbing PCF p forms an equilibrium if the following two conditions hold, for any $(h, x) \in H$:

1. If $p(h, x) \in X\backslash\{x\}$, then $p(h, x)$ is an efficient and weakly preferred move from x for $S(x, p(h, x))$.

2. If $p(h, x) = x$, then x is an efficient move from x for all S.

Given this equilibrium criterion, it is now clear that this version of dynamic equilibrium process of coalition formation is equivalent with the one defined for terminating processes. That is, the set of states that are absorbing after some history in some equilibrium of the KR model of real time blocking coincides with the set of outcomes that are implementable after some history in a corresponding equilibrium of the one-shot game.

However, as hinted above, our solution differs from KR in two ways (apart from allowing the process to be history dependent). First, part 2 of the condition in KR is slightly stronger as it requires that if the status quo state x is not efficient for some coalition, then the equilibrium move must be efficient and *strictly* preferred move from x for some coalition (we only demand weak preference, by part 1 of the condition).[18] This additional desideratum would affect the sufficiency but not

[18]The motivation for this weakening is our focus on properties of the solution that stem from the one-deviation principle.

the necessity of the characterization in Theorem 1. Second, in the original formulation of KR, intertemporal preferences are captured by the discounting criterion. Formally, given a discount factor $\delta \in (0,1)$, player i's value function V_i^δ is derived as the solution to the recursive expression (see Ray, 2007)

$$V_i^\delta(h,x) = (1-\delta)u_i(x) + \delta V_i^\delta(h,x,p(h,x)). \tag{11}$$

KR define the coalitional preferences and efficiency criteria and, *a fortiori*, the solution vis-a-vís the these value functions.[19]

We should add that there are games in which the way intertemporal payoffs are evaluated (even in the limit) does make a difference. An example is given in Appendix A.

5.2 Largest Consistent Set

Rosenthal (1972) and Chwe (1994) define a general framework of coalitiuon formation used in this work. Chwe's (1994) aim is to rule out outcomes whose stability could never be an issue. His solution, the largest consistent set, always exists and is one of the most frequently used solutions in applications. We demonstrate that the set of outcomes that are implementable via any dynamic equilibrium process of coalition formation is a subset (sometimes strict) of the largest consistent set.

The now we develop Chwe's (1994) solution. To this end, an outcome $y \in X$ *indirectly* dominates x if there is a feasible path $(x_0, ..., x_K) \in \mathcal{F}$ such that $x_0 = x$, $y = x_K$ (recall our the definition (2) of a feasible path).[20] Set $C \subseteq X$ is a *consistent set* if C consists of all x for which the following holds: if $z \in F_S(x)$, then there is $y \in C$ such that either $y = z$ or y indirectly dominates z and $y \not\succ_S x$. Chwe showed that a consistent set exists and the largest (in the sense of set inclusion) consistent set is unique.

We show that, in terms of implementable outcomes, dynamic equilibrium process of coalition formation is a refinement of the largest consistent set. We do this by first showing that the ultimate uncovered set implements a consistent set of outcomes.

Proposition 1 *$\mu[UUC]$ is a subset of the largest consistent set.*

Proof. It suffices to show that $[UUC]$ is a consistent set. Suppose, on the contrary, that this is not the case. Then there is an $x \in \mu[UUC]$, an S, and a $z \in F_S(x)$

[19] In KR, a history dependent version of weak preference reads $V_i^\delta(h,x,y) \geq V_i^\delta(h,x)$ for all i and for all S. However, given the definition of V_i^δ, this is equivalent to $V_i^\delta(h,x,y) \geq u_i(x)$ for all i and for all S.

[20] As Konishi and Ray (2003) and Ray (2007), we appeal to the weaker version of Chwe's (1994) notion of indirect dominance. The original one requires strict preference.

such that $\mu[\bar{y}] \succ_S x$, for all $\bar{y} \in \mathcal{F}_z \cap UUC$. But since there is $\bar{x} \in UUC$ such that $\mu[\bar{x}] = x$, this contradicts the assumption that $UUC = uc(UUC)$. That is, that $\bar{x}$ is not covered in UUC. ■

Importantly, $\mu[UUC]$ need not coincide with the largest consistent set. We now show via an example that $\mu[UUC]$ can be a *strict* subset of the largest consistent set. Consider the game depicted in Figure 2.

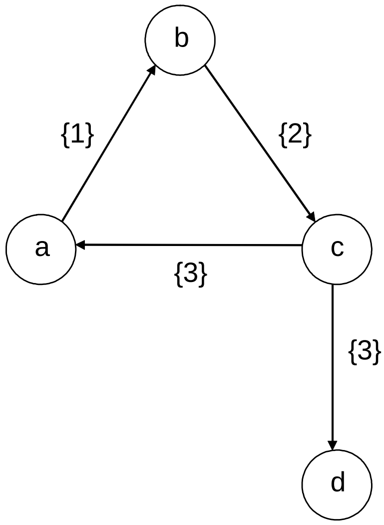

Figure 2

Here $N = \{1, 2, 3\}$, $X = \{a, b, c, d\}$. Numerical payoffs (in the order of player indices) from choices a, b, c, and d are, respectively, $(0, 0, 1)$, $(0, 1, 0)$, $(1, 0, 0)$, and $(2, 2, 2)$. The largest consistent set is $\{a, b, d\}$. However, the largest consistent path structure UUC consists the paths $\{(a, b, c, d), (b, c, d), (c, d), (d)\}$, and hence $\mu[UUC] = \{d\}$.

KR show that all absorbing states of stationary absorbing deterministic equilibrum processes that have the Markovian property lie in the largest consistent set. Our results thus indicate that this property holds even without the Markovian restriction.

5.3 Stable Standard of Behavior and a Full Coalitional Equilibrium

Mariotti (1997) and Xue (1998) are motivated by considerations that are closely related to this paper (see also Mariotti and Xue, 2003). In this section, we discuss their solutions, *full coalitional equilibrium* and *stable standard of behavior*, respectively, from the viewpoint of dynamic equilibrium processes of coalition formation.

Xue (1998), arguing that the largest consistent set is too permissive, imposes a consistency criterion on paths rather than on outcomes. His starting point is the observation that - since in a dynamic environment play can move *in equilibrium* from one status quo to another before an outcome is implemented - what is crucial is the stability of paths rather than that of the outcomes. Xue (1998) uses Greenberg's (1991) foundational *Theory of Social Situations* to identify paths that are robust against deviations. This theory is based on the von Neumann-Morgenstern stable set approach. This solution concept bears similarity to the consistent path structure, the reduced form characterization of our solution. The difference between Xue's (1998) solution and ours is that the former permits deviations also by inactive coalitions. This property has concequences on the existence of the solution.

To interpret our results in the framework of Xue (1998), let us lay down the key ingredients of his model.[21] Assume that the alternative $x \in X$ is the current status quo. Consider a path $\bar{x}$ and some of its node x_k. Assume that a coalition S can replace x_k by some alternative $y \neq x_{k+1}$. In doing so, S is aware of that the set of feasible paths from y is $\mathcal{F}_y$. In contemplating such a deviation from y, however, members of S base their decision on comparing paths that might be followed by rational and farsighted individuals at y. Let $SB(y) \subset \mathcal{F}_y$ denote this *standard of behavior*.

The following definition, which is adopted from Xue (1998), describes in our notation the conservative approach to stable standard of behavior.

Definition 5 *A standard of behavior SB is* conservatively stable *if it is*

Internally stable*: for all $x \in X$, if $\bar{x} \in SB(x)$, then there is* no *y, coalition S, and k such that $y \in F_S(x_k) \setminus \{x_{k+1}\}$, and $\mu[\bar{y}] \succ_S \mu[\bar{x}]$, for all $\bar{y} \in SB(y)$,*

Externally stable*: for all $x \in X$, if $\bar{x} \in \mathcal{F}_x \setminus SB(x)$, then there is y, coalition S, and k such that $y \in F_S(x_k) \setminus \{x_{k+1}\}$, and $\mu[\bar{y}] \succ_S \mu[\bar{x}]$, for all $\bar{y} \in SB(y)$.*

To see most transparently the relationship between our solution concept and the conservatively stable standard of behavior of Xue (1998), let us spell out the key features of the feasible ultimate uncovered set in the language of internal and external stability.

Remark 1 *The ultimate uncovered set UUC satisfies*

Internal stability*: if $\bar{x} \in UUC$, then there is* no *y and k such that $y \in F_S(x_k) \setminus \{x_{k+1}\}$ and $\mu[\bar{y}] \succ_S \mu[\bar{x}]$, for $S = S(x_{k+1}, x_k)$, and for all $\bar{y} \in \mathcal{F}_y \cap UUC$.*

External stability*: if $\bar{x} \in \mathcal{F} \setminus UUC$, then there is y and k such that $y \in F_S(x_k) \setminus \{x_{k+1}\}$ and $\mu[\bar{y}] \succ_S \mu[\bar{x}]$, for $S = S(x_{k+1}, x_k)$, and for all $\bar{y} \in \mathcal{F}_y \cap UUC$.*

[21] Referring to Greenberg's (1991) Theory of Social Situations.

Internal stability follows from the definition of consistent path structure. External stability, in turn, is a consequence of the construction of UUC; every path outside UUC is covered in UUC. It is easy to see that the key difference between the solution concepts is that the notion of stable standard of behavior requires stability against arbitrary coalitional deviations whereas the ultimate uncovered set only with respect to the active coalition. Since the restriction affects both internal and external stability but to the different directions, there is no straightforward relationship between the concepts.

From these definitions it is clear that the notions of stable standard of behavior and the feasible ultimate uncovered set coincide in games where the collection of possibly active coalitions is singleton, *i.e.,* where only one coalition may change the status quo at a time. Roughly, this corresponds to the scenarios where the coalitional game can be exhibited in the form of a simple recursive game (vis-a-vis interpreting coalitional preferences as a single individual's preferences).

The more stringent requirement for acceptable deviations prevent pathological blocking relationships and guarantees the existence of an ultimate uncovered set. Conversely, the leeway provided by arbitrary coalitional deviations in the context of stable standard of behavior is the reason for the occasional emptiness of the solution, as is the case in the following example (due to Xue, 1998):

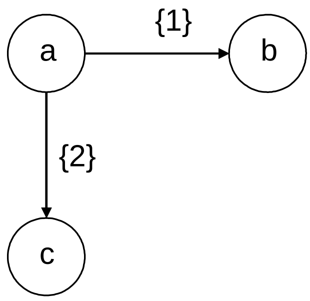

Figure 3

where payoffs from a, b and c are, respectively, $(0,0)$, $(2,1)$, and $(1,2)$. The unique conservative standard of behavior is empty. However, the ultimate uncovered set contains both the arcs: $\mathcal{F}_a \cap UUC = \{(a,b),(a,c)\}$ and hence $\mu[UUC] = \{b,c\}$.

To the best of our knowledge, Mariotti (1997) is the only paper in the literature that explicitly models history dependent coalitional strategies. Although his set up is different from this paper, there are notable similarities in how strategic behavior is modeled. In fact, we argue that all conditions imposed on his full coalitional equilibrium -solution are also satisfied by our solution. Hence, our existence result can also be applied to Mariotti's set up.

Let us sketch Mariotti's (1997) model, where players seek to form an agreement of how to play a strategic form game. The game is defined by sets A_i of actions for each player i. Preferences are defined over the set of all profiles of possible actions $A_N = \times_{i \in N} A_i$. The play proceeds by coalitions S proposing a profile of actions in $A_S = \times_{i \in S} A_i$. If coalition S serves a the proposer, then the choices of the complementary coalition $-S$ are herited from the current status quo profile. Hence, the set of feasible outcomes of coalition S, given the status quo profile $x \in X := A_N$, is obtained by

$$F_S(x) = \{(y_S, x_{-S}) : y_S \in A_S\} \subseteq X.$$

Now we can apply Definition 1 to specify conditions under which an agreement in reached.

The key features of Mariotti's (1997) novel approach bear similarity to ours: The strategies are assumed history dependent and the focus is, at the outset, on strategies that implement an outcome in finite time. Also a version of the one-deviation property is used as the equilibrium criterion. While Mariotti's solution, full coalitional equilibrium, is richer in a sense that it allows the coalitions to disagree on how the play will evolve in the future, it is apparent the solution *also* permits the agents to agree on how the play proceeds. This is our basic assumption: the strategy specifies a unique future play path that is known by all players. Therefore we conclude that, in the case of a strategic form games, a dynamic equilibrium process of coalition formation is a particular case of a full coalitional equilibrium of Mariotti (1997). By Corollary 2 it then follows that also *a full coalitional equilibrium exists.*[22] However, of there may well be many full coalitional equilibria of which our characterization cannot say much.

6 Efficiency

A recurrent theme in coalitional analysis is efficiency. A classic argument that goes under the name of the Coase theorem says that an outcome that results from unrestricted bargaining will always be efficient: otherwise it would be blocked by another outcome that all players prefer. This intuition is not sufficient in the current framwork.

Rosenthal's centipede game is an example of a game in which the dynamic equilibrium process of coalition formation does not implement an efficient outcome (in finite games of perfect information, the equilibrium coincides with the backwards induction -solution). But the centipede game is not a good example since the Coase theorem appeals to the coalitions' inability to *commit* to a bad outcome. This implicitly assumes that coalitional moves do not have irrecoverable

[22]There are minor differences in the defining inequalities of the two solutions concepts. However, such differences do not affect the qualitative nature of the results

consequences. This desideratum is met for example by games that are *irreducible* in a sense that from any node there is a path to any other node.

It turns out that irreducibility is not sufficient for efficiency. Consider the following anarchistic economy. There are two players 1 and 2. Assume that initial utility allocation is unequal, say 4 to player 1, and 1 to player 2. The worse off player, in this case 2, can pillage utility units from the other player, in this case from player 1 (cf. Jordan, 2005). Pillaging is costly so that 2 receives only one half of what he pillages. After pillaging, the players can increase, by common effort, both of their payoffs by one unit. This results in payoffs 1 and 4 for players 1 and 2, respectively.[23] Then the roles of the players are reversed in a sense that now the worse off player is 1 and he may pillage player 2's utility units. The game is depicted in Figure 4 where the directed edges reflect potential coalitional moves (numbers inside the nodes reflects the payoffs). The graph constitutes a circuit. Nodes $(3, 0)$ and $(0, 3)$ are Pareto dominated by the next nodes in the circuit. Nevertheless the largest consistent set, and hence $\mu[UUC]$ by Proposition 1, consists only of the low payoff nodes $(3, 0)$ and $(0, 3)$. Thus *every* dynamic equilibrium process of coalition formation implements only Pareto dominated outcomes.

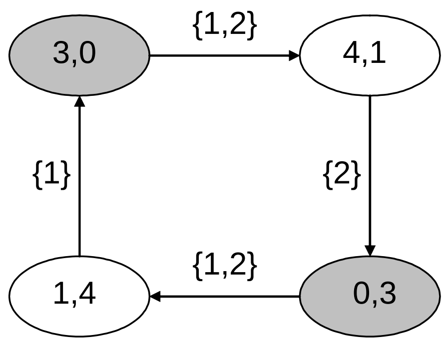

Figure 4

One can argue that the inefficiency here is due to artificially restricted blocking relationships. A stronger requirement than irreducibility is *completeness*: for any y and x there is S such that $y \in F_S(x)$. But again, there are irreducible and complete games where $\mu[UUC]$ contains *only* Pareto-dominated outcomes. An example is depicted in Appendix B.

A third, yet stronger condition on the graph would be that $F_S(x) = X$ *for all* S, so that at each node all coalitions can move the play to any node. In such case

[23] Pillaging by 2 is not possible when payoffs are $(3, 0)$ since pillaging itself requires some resources.

there would be no *a priori* restrictions on what a coalition can achieve. It is easy to show that this restriction would not preclude any Pareto-optimal outcome from being in $\mu[UUC]$. However, Pareto-optimality is still not guaranteed. Modify, for example the game depicted in Figure 4 by allowing each coalition $\{1\}, \{2\}$, and $\{1, 2\}$ induce any outcome from any outcome. Then $\mu[UUC]$ would contain all the outcomes but there would also be a consistent path structure $\mathcal{C}$ such that $\mu[\mathcal{C}] = \{(3, 0), (0, 3)\}$. Hence the correponding equiliubrium process could only implement Pareto dominated outcomes. We therefore have to conclude that there is nothing inconsistent in the idea that an inefficient outcome becomes implemented, even if bargaining opportunities are unrestricted.[24]

7 Discussion

We have showed that a dynamic version of the absorbing, deterministic equilibrium process of coalition formation of Konishi and Ray (2003) always exists. The crucial assumption is that the process may now depend on history of the play. Another important assumption is that only one coalition may be active in blocking the underlying status quo. This assumption may, of course, be criticized.

Consider a following example due to Xue (1998):

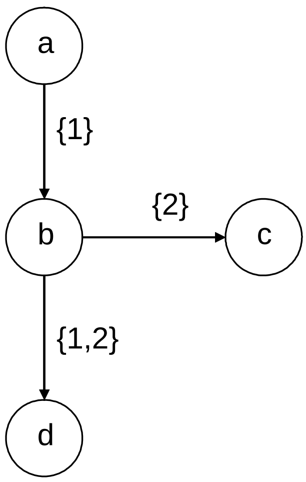

Figure 5

Now $N = \{1, 2\}$, $X = \{a, b, c, d\}$, and numerical payoffs (in the order of players' indices) from a, b, c, and d are, respectively, $(6, 0), (7, 4), (5, 10), (10, 5)$. The largest

[24] Konishi and Ray (2003) establish conditions on strategies under which Pareto-optimality is guaranteed.

consistent picks the same set of outcomes $\{a, c, d\}$ that is implemented by the dynamic equilibrium processes of coalition formation. However, as Xue (1998) argues, this set is too large since d should not ever be chosen: when a is the status quo, the "predicted" outcomes are $\{a, d\}$, the latter when the coalition $\{1, 2\}$ induces d. But once the coalition $\{1, 2\}$ is about to form, player 2 would renege the contract and choose the option c instead. Hence, d should not be considered as a good prediction.

This is a valid argument in many circumstances. A way to cope with the criticism is to require robustness against *nested* coalitional deviations.[25] Then part 1 of Definition 1 would read,

1'. $\sigma(h, x) \in X$ implies $\sigma(h, x)$ is an efficient and weakly preferred move from x for *all* coalitions $S \subseteq S(x, y)$.

To recover the key results with nested deviations, one would only need to define a new path dominance operator, $\rhd^*$, this time with respect to all subsets of the active coalition:

A path $\bar{y} \in \mathcal{F}$ $*-$dominates a path $\bar{x} \in \mathcal{X}$ at the k^{th} step, denoted by $\bar{y} \rhd_k^* \bar{x}$, if $\mu[\bar{y}] \succ_S \mu[\bar{x}]$, for all $\bar{y} \in \mathcal{F}_y$, for some $y \in F_S(x_k)$, for some $S \subseteq S(x_k, x_{k+1})$.

By replacing "$\rhd$ " with "$\rhd^*$ " in the definition of a consistent path structure and the property "$S = S(x_k, x_{k+1})$" with "for any $S \subseteq S(x_k, x_{k+1})$" in the proofs, the existence of a solution would be obtained along the same avenues as in the current version. The new solution would again implement a set of outcomes that is a refinement of the largest consistent set.

But even if one wants to incorporate the idea of nested deviations into the framework, it is not clear why one should do it through the solution concept. An alternative - and perhaps cleaner - way to model this is to assume that the game form Γ itself reflects nested deviations. This would be more in line with the standard position in game theory that all (relevant) behavioral *opportunities* should be specified through the game form (and not by the solution concept).[26] In the coalitional set up, this principle implies that the solution concept should only describe optimal behavior of a single behavioral unit, a coalition.

More formally, to amend nested deviations into the game form, think of forming a derived game Γ' from the original one Γ that has two kinds of nodes, physical X (as in the game Γ) and transitional Z (in the present framework, transitional

[25] Some nestedness assumption is often used to guarantee the existence. The classic papers in the field, Hart and Kurz 1983, Bernheim et al. 1987, and Ray and Vohra 1997, are cases in point.

[26] See e.g. Osborne and Rubinstein (1994).

nodes could be modeled as "bad" outcomes). The derived game Γ' would now have the property that if $y \in F_S(x)$, then there are a transitional nodes $z_0, ..., z_{|S|+1}$ such that $z_0 = x$ and $z_{|S|+1} = y$, for some indexing j of the subsets of S, and such that $S^j \subset S$ would be the only potentially active coalition under z_j with $F'_{S^j}(z_j) = F_{S'}(x) \cup \{z_{j+1}\}$, for all $j = 0, ..., |S|$. Then the derived game would offer each subset of the original game's active coalition a chance to redirect the play to the direction it prefers, if it does not like the choice made by the originally active coalition. By nestedness of S^j and S, such derived game could clearly be constructed. In a dynamic equilibrium process of coalition formation, this would guarantee that no subset of the original active coalition would not want to deviate further.

References

[1] Ambrus A. (2006), Coalitional rationalizability, Quarterly Journal of Economics 121, 903-29

[2] Aumann, R. and R. Myerson (1988), Endogenous formation of links between players and of coalitions: an application of the Shapley value, in: A. Roth (ed.), The Shapley Value: Essays in Honor of Lloyd Shapley, Cambridge University Press, Cambridge, UK, 175–91

[3] Barberà, S. and A. Gerber, (2003), On coalition formation: durable coalition structures, Mathematical Social Sciences 45, 185–203

[4] Barberà, S. and A. Gerber (2007), A note on the impossibility of a satisfactory concept of stability for coalition formation games, Economics Letters 95, 85-90

[5] Bernheim, D., B. Peleg, and M. Whinston (1987), Coalition-proof Nash equilibria I, Concepts, Journal of Economic Theory 42, 1-12

[6] Bloch, F. (1996), Sequential Formation of Coalitions in Games with Externalities and Fixed Payoff Division, Games and Economic Behavior 14, 90-123

[7] Bloch, F. and A. Gomes (2006), Contracting with externalities and outside options, Journal of Economic Theory 127, 172-201

[8] Chatterjee, K., B. Dutta., D. Ray, and K. Sengupta. (1993), A noncooperative theory of coalitional bargaining, Review of Economic Studies 60, 463-77

[9] Chwe, M. (1994), Farsighted stability, Journal of Economic Theory 63, 299-325

[10] Conley, J. P. and H. Konishi (2002), Migration-Proof Tiebout Equilibrium: Existence and Asymptotic Efficiency, Journal of Public Economics 86, 243-62

[11] Dutta, B. and S. Mutuswami (1997), Stable networks, Journal of Economic Theory 76, 322–344

[12] Diamantoudi, E. and L. Xue (2006), Coalitions, agreements and efficiency, forthcoming in Journal of Economic Theory

[13] Dutta. B (1988), Covering sets and new condorcet correpondence, Journal of Economic Theory 44, 63-80

[14] Fishburn, P. (1977), Condorcet social choice function, SIAM Journal of Applied Math 33, 295-306

[15] Greenberg, J.(1990), The theory of social situations, Cambridge University Press, Cambridge UK

[16] Gomez, A. and P. Jehiel (2005). Dynamic process of social and economic interactions: on the persistence of inefficiencies, Journal of Political Economy 113, 626-67

[17] Hart, S. and M. Kurz (1983), Endogenous formation of coalitions, Econometrica 51, 1047–64

[18] Herings, J.-J., Mauleon, A and V. Vannetelbosch (2004), Rationalizability for social environments, Games and Economic Behavior 49: 1, 135-56

[19] Harsanyi, J. (1974), An equilibrium point interpretation of stable sets and a proposed alternative definition, Management Science 20, 1427-9

[20] Hyndman, K. and D. Ray, (2007), Coalition Formation with Binding Agreements, Review of Economic Studies 74, 1125-47

[21] Jackson, M. (2008), Social and Economic Networks, Princeton University Press, Princeton NJ

[22] Jackson, M. and A. van den Nouweland (2005) Strongly Stable Networks, Games and Economic Behavior 51, 420-444

[23] Jackson, M. and A. Wolinsky (1996), A strategic model of social and economic networks, Journal of Economic Theory 71, 44–74

[24] Konishi, H. and D. Ray (2003), Coalition formation as a dynamic process, Journal of Economic Theory 110, 1-41

[25] Flesch, J., Kuipers, J., Schoenmakers, G. and K. Vrieze, (2008), Subgame-Perfection in Stochastic Games with Perfect Information and Recursive Payoffs, METEOR Research Memoranda 041

[26] Laslier, J.-F. (1997), Tournament Solutions and Majority Voting, Heidelberg, New-York: Springer-Verlag

[27] Mariotti, M. (1997), A model of agreements in strategic form games, Journal of Economic Theory 74, 196–217

[28] Mariotti, M. and L. Xue (2003), Farsightedness in coalition formation, in Carlo Carraro (ed.) The Endogenous Formation of Economic Coalitions, Edward Elgar, 128-55

[29] Mauleon, A.and V. Vannetelbosch (2004), Farsightedness and Cautiousness in Coalition Formation Games with Positive Spillovers, Theory and Decision 56, 291-324

[30] Miller, N. (1980), A new solution set to for tournaments and majority voting, American Journal of Political Science 24, 68-96

[31] Page, F., Kamut, S. and M. Wooders (2005), Networks and farsighted stability, Journal of Economic Theory 120, 257-61

[32] Page, F. and M. Wooders (2006), Networks and Clubs, working paper, University of Warwick

[33] Osborne, M. and A. Rubinstein (1994), A course in game theory, MIT Press, Cambridge MA

[34] Ray, D. (2007), A game theoretic perspective on coalition formation, The Lipsey Lectures, Oxford University Press, Oxford UK

[35] Ray, D. and R. Vohra (1997), Equilibrium binding agreements, Journal of Economic Theory 73, 30-78

[36] Ray, D. and R. Vohra (1999), A theory of endogenous coalition structures, Games and Economic Behavior 26, 286-336

[37] Rosenthal, R. (1972), Cooperative games in effectiveness form, Journal Economic Theory 5, 88–101

[38] Xue, L. (1997), Nonemptiness of the largest consistent set , Journal of Economic Theory 73, 453-59

[39] Xue, L (1998), Coalitional stability under perfect foresight, Economic Theory 11, 603-27

[40] Vartiainen, H. (2005), Cognitive equilibrium, ACE working paper 20

[41] von Neumann, J. and O. Morgenstern (1947), Theory of games and economic pehavior, Princeton University Press, Princeton NJ

[42] Yi, S.-S. (1997), Stable coalition structures with externalities, Games and Economic Behavior 20, 201–237

A Appendix

Consider the game depicted in Figure 6.

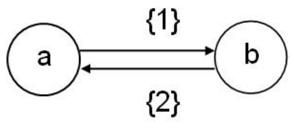

Figure 6

Here $N = \{1, 2\}$, $X = \{a, b\}$, and $x \xrightarrow{S} y$ denotes the relation $y \in F_S(x)$ for $y \neq x$. Payoffs (in the order of player indices) from the states a and b are, respectively, $(0, 1)$ and $(1, 0)$. Under discounting, there is no determinstic, absorbing equilibrium process of coalition formation. However, as Corollary 2 argues, under limit-of-the-means criterion such equilibrium PCF does exist.

B Appendix

Consider the three players/six alternatives game, depicted in Figure 7. The directed graph describes the possible coalitional blockings between the nodes. The graph is complete but, for simplicity, we only depict blockings to one direction, *i.e.*, $(3, 1, 2) \rightarrow_{\{2\}} (1, 2, 0)$ means that also $(1, 2, 0) \rightarrow_{\{1,3\}} (3, 1, 2)$, etc.. Again, $\mu[UUC]$ consists of the shaded nodes as any Pareto-optimal, non-shaded node is covered via a shaded node.

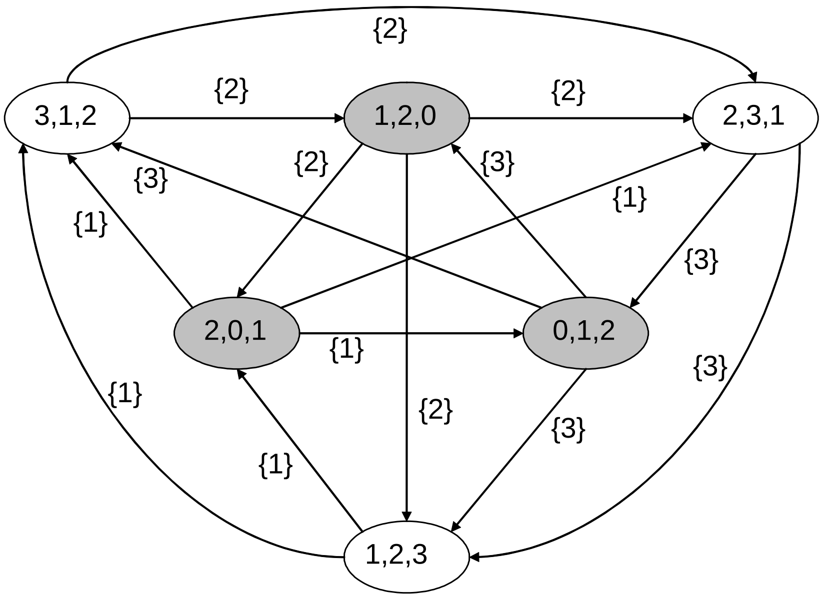

Figure 7