Collusion, Randomization, and Leadership in Groups☆

Rohan Dutta[1], David K. Levine[2], Salvatore Modica[3]

Abstract

We propose a theory of collusive groups in the context of finite non-cooperative games. We first consider a simple setting in which players are exogenously partitioned into groups within which players are symmetric. Given the play of the other groups there may be several symmetric equilibria for a particular group. We develop the idea that if a group can collude they will agree to choose the equilibrium most favorable for its members. This leads to an equilibrium concept which we call *collusion constrained equilibrium*. We then consider an alternative model of a non-cooperative meta-game played between leaders of groups who issue instructions and evaluators who carry out *ex post* punishment of the leaders if the instructions fail to be incentive compatible. We establish equivalence between equilibria of the collusive group game and the leadership game.

We extend the leadership model to games where players within groups are not necessarily symmetric and groups are endogenously formed. In this model leaders compete for followers by making credible offers of the utility followers will receive if they play according to the leader's instructions. This leads to a rich theory of group formation which we explore through a series of examples. We find robust equilibria that involve mixing and Pareto superior equilibria that do not involve mixing but are less robust to the leadership structure. We also show in prisoners' dilemma type settings that the frequency of cooperation increases as the benefit to cooperation increases and the benefit of deviating decreases.

Keywords: Collusion, Coalition Formation, Organization, Group

☆First Version: June 2, 2013. We are especially grateful to Michele Boldrin, Michael Chwe, Drew Fudenberg, and Zacharias Maniadis. We are grateful to the EIEF, to NSF Grant SES-08-51315 and to the MIUR PRIN 20103S5RN3 for financial support.

*Corresponding author David K. Levine, 1 Brooking Dr., St. Louis, MO, USA 63130

Email addresses: `rohan.dutta@mcgill.ca` (Rohan Dutta), `david@dklevine.com` (David K. Levine), `salvatore.modica@unipa.it` (Salvatore Modica)

[1]Department of Economics, McGill University

[2]Department of Economics, EUI and WUSTL

[3]Università di Palermo

1. Introduction

Individuals often act as members of a group, but groups do not act as individuals, that is although they tend to act collusively, groups or their leaders cannot ignore members' individual preferences and incentives. Hence neither individual rationality alone nor "group rationality" alone are suited to analyze strategic interaction between individuals mediated by groups. We propose a theory of collusive groups in the context of finite non-cooperative games. We first consider a simple setting in which players are exogenously partitioned into groups within which players are symmetric. Given the play of the other groups there may be several symmetric Nash equilibria within a particular group. The idea is that a collusive group will agree to choose the equilibrium most favorable for its members. This leads to an existence problem, which we illustrate with an example. We overcome the problem through a type of randomization that eliminates a discontinuity, leading to what we call *collusion constrained equilibrium*. We show how these equilibria arise as limits of belief perturbed models in which groups do collude on the equilibrium most favorable to its members. We then consider an alternative model of a strictly non-cooperative meta-game played between "leaders" and "evaluators" of groups. We establish equivalence between equilibria of the collusive group game and the leadership game.

We then extend the leadership model to games where players within groups are not necessarily symmetric and groups are endogenously determined by competition of leaders. In this model leaders bid for followers by making credible offers of the utility followers will receive if they follow their instructions. This leads to a rich theory of group formation which we explore through a series of examples. We find robust equilibria that involve mixing and Pareto superior equilibria that do not involve mixing but are less robust to the leadership structure. We also show in prisoners' dilemma type settings that the frequency of cooperation increases as the benefit to cooperation increases and the benefit of deviating decreases. There are two basic questions that we address. First we ask the extent to which particular leadership structures can improve outcome efficiency. We show how this can be the case and how it depends on leadership structure. Second we ask how institutions impact on leadership structure. This we address by examining how the leadership structure matters and the constraints that determine leadership structure.

The issue that groups do not behave as a single individual has been discussed as well in the literature on collective action (for example Olson (1965)), but that literature has not provided a general framework for analysis, proposing instead particular solutions such as tying arrangements or other commitments to overcome incentive constraints. The branch of the game theory literature that is most closely connected to the ideas we develop here is the literature that uses non-cooperative methods to analyze cooperative games and in particular the endogenous formation of coalitions. One example is Ray and Vohra (1999) who introduce a game in which players bargain over the formation of coalitions by making proposals to coalitions and accepting or rejecting those proposals within coalitions. This literature generally describes the game by means of a characteristic function and involves proposals and bargaining. Although our model of endogenous group formation also involves an element of bidding, we work in a framework of implicit or explicit coordination among group members in a non-cooperative game among groups. This is similar in spirit to Bernheim, Peleg and Whinston (1987)'s variation on strong Nash equilibrium, that they call coalition-proof Nash equilibrium, although the details of our model are rather different.

There is a long literature on collusion in mechanism design, and our model builds on those ideas. With a few exceptions the general idea is that within a mechanism a particular group - the bidders in an auction, the supervisor and agent in the Principal/Supervisor/Agent model, for example - must not wish to recontract in an incentive compatible way. In the case of the hierarchical models, the Principal/Agent/Supervisor model of collusion originates with Tirole (1986) and the more general literature on hierarchical models is discussed in his survey Tirole (1992); for a recent contribution and an indication of the current state of the literature see Celik (2009). In the auction literature, we mention the papers of McAfee and McMillan (1992) and Caillaud and Jéhiel (1998) among others. The theory has been pursued for other types of mechanisms, as in Laffont and Martimort (1997). In most of this work there is only one group recontracting, so the issue of a "game" among the groups does not arise. Our setting involves multiple groups on an equal footing. The closest model we know of is that of Che and Kim (2009) in the auction setting - they allow multiple groups they refer to as cartels to recontract in an incentive compatible way among themselves. However, it does not appear that strictly speaking these cartels play a game. Similar to the mechanism design literature is the study of collusion in monetary matching models such as in Hu, Kennan and Wallace (2009) where pairs of players who are matched can choose their most preferred equilibrium within the pair. It is also the case that in the theory of clubs, such as Cole and Prescott (1997) and Ellickson et al (2001), implicitly collusion takes place within clubs - although the clubs interact in a market rather than game environment.

In applied work - for example by economic historians - the issue of how groups behave is usually dodged by examining a game in which an entire group is treated as a single individual. This is the case in the current literature on the role of taxation by the monarchy in bringing about more democratic institutions. Hoffman and Rosenthal (2000) explicitly assume that the monarch and the elite act as single agents, and this assumption seems to be accepted by later writers such as Dincecco, Federico and Vindigni (2011). As the literature on collusion in mechanism design makes clear, by treating a group as an individual we ignore the fact that the group itself is subject to incentive constraints. Individuals wish other individuals to act in the group interest, but may not wish to do so themselves. In a sense we generalize the literature that assumes that a group decision is made by a single leader by adding to the game an evaluator for that leader who punishes the leader for violating incentive constraints.

Leadership is also studied in models where a group benefits from its members coordinating their actions in the presence of imperfect information about the environment - see for example Hermalin (1998), Dewan and Myatt (2008) and Bolton, Brunnermeier and Veldkamp (2013). As in the present paper the leader provides guide for action, but the similarity ends more or less there. In these papers there is no game between groups (the focus of our paper), the problem is how to exploit the information being acquired by leader and group members in the group interest, and what in particular are desirable properties of the leader's decision process. To be concrete, Bolton, Brunnermeier and Veldkamp (2013) find for example that the leader should not put too much weight on the information coming from followers (what they call "resoluteness" of the leader). We focus on strategic interaction between groups, so our model of interaction between leader and group members is much coarser than in the cited papers: the leader proposes a common course of action and all group members take the same action in equilibrium. On the other hand, a central element of our models is accountability, in that a leader whose recommendations are not endorsed by the group will be punished.

2. A Motivating Example

The simplest - and as indicated in the introduction a widely used - theory of collusion is one in which players are exogenously divided into homogeneous groups subject to incentive constraints. If - given the play of other groups - there is more than one in-group symmetric equilibrium then a group should be able to agree or coordinate on their "most desired" equilibrium.

Example 1. We start with an example with three players. The first two players form a collusive group while the third acts independently. The obvious condition to impose in this setting is that given the play of player 3, players 1 and 2 should agree on the incentive compatible (mixed) common action that gives them the most utility. However, in the following game there is no equilibrium that satisfies this prescription. Specifically, each player chooses one of two actions, C or D and the payoffs can be written in bi-matrix form. If player 3 plays C the payoff matrix for the actions of players 1 and 2 is a symmetric Prisoner's Dilemma game in which player 3 prefers that 1 and 2 cooperate (play C)

	C	D
C	$6,6,5$	$0,8,5$
D	$8,0,5$	$2,2,0$

If player 3 plays D the payoff matrix for the actions of players 1 and 2 is a symmetric coordination game in which player 3 prefers that 1 and 2 defect (play D)

	C	D
C	$6,6,0$	$4,4,0$
D	$4,4,0$	$5,5,5$

Let α^i denote the probability with which player i plays C. We examine the set of equilibria for players 1 and 2 given the strategy α^3 of player 3. If $\alpha^3 > 1/2$ then D is strictly dominant for both player 1 and 2 so there is a unique in-group equilibrium in which they play D, D. If $\alpha^3 = 1/2$ then there are two equilibria, both symmetric, one at C, C and one at D, D. If $\alpha^3 < 1/2$ then there are three equilibria, all symmetric, one at C, C, one at D, D and a strictly mixed equilibrium in which $\alpha^1 = \alpha^2 = (1/3)(1+\alpha^3)/(1-\alpha^3)$.[4]

How should the group of player 1 and player 2 collude given the play of player 3? If $\alpha^3 > 1/2$ they have no choice: there is only one in-group equilibrium at D, D. For $\alpha^3 \leq 1/2$ they each get 6 at the C, C equilibrium, no more than 5 at the D, D equilibrium, and strictly less than 6 at the strictly mixed equilibrium. So if $\alpha^3 \leq 1/2$ they should choose C, C. Notice that in this example there is no ambiguity about the preferences of the group: they unanimously agree in each case as to which is the best equilibrium.

We may summarize the play of the group by the "group best response". If $\alpha^3 > 1/2$ then the group plays D, D while if $\alpha^3 \leq 1/2$ the group plays C, C. What is the best response of player 3 to the play of the group? When the group plays D, D player 3 should play D and so $\alpha^3 = 0$ and in particular is not larger than $1/2$; when the group plays C, C player 3 should play C and so $\alpha^3 = 1$ and in particular is not less than or equal to $1/2$. Hence there is no equilibrium of the game in which the group of player 1 and player 2 chooses the best in-group equilibrium given the play of player 3.

In this example, the non-existence of an equilibrium in which player 1 and player 2 collude is driven by the discontinuity in the group best response: a small change in the probability of α^3 leads to an abrupt change in the behavior of the group. The key idea of this paper is that this discontinuity is an artifact of the model and does not make sense from an economic point of view. In particular, it does not make much sense that as α^3 is increased

[4]The structure of group equilibria is easily seen by considering the payoff to player 1 if he plays C or D:

$$C: \quad \alpha^2 \cdot [\alpha^3 \cdot 6 + (1-\alpha^3)6] + (1-\alpha^2)[\alpha^3 \cdot 0 + (1-\alpha^3)4]$$
$$D: \quad \alpha^2[\alpha^3 \cdot 8 + (1-\alpha^3)4] + (1-\alpha^2)[\alpha^3 \cdot 2 + (1-\alpha^3)5]$$

From this it is clear that for the group: D, D is equilibrium for all α^3, strictly incentive compatible; If $\alpha^3 > 1/2$ then D, D is the unique equilibrium; If $\alpha^3 = 1/2$ then also C, C is equilibrium - weakly incentive compatible; If $\alpha^3 < 1/2$ then C, C is also equilibrium, and it is strictly incentive compatible. For the mixed equilibrium, the condition that player 1 must be indifferent between C and D gives the following:

$$6\alpha^2 + (1-\alpha^2)4(1-\alpha^3) = \alpha^2(8\alpha^3 + 4(1-\alpha^3)) + (1-\alpha^2)(2\alpha^3 + 5(1-\alpha^3))$$
$$\left(6 - 4(1-\alpha^3)\right)\alpha^2 + 4(1-\alpha^3) = \left(8\alpha^3 + 4(1-\alpha^3) - 2\alpha^3 - 5(1-\alpha^3)\right)\alpha^2 + (2\alpha^3 + 5(1-\alpha^3))$$
$$\left(6 - 4(1-\alpha^3)\right)\alpha^2 + 4 - 4\alpha^3 = \alpha^2(6\alpha^3 - (1-\alpha^3)) + 5 - 3\alpha^3$$
$$\left(6 - 3(1-\alpha^3) - 6\alpha^3\right)\alpha^2 = 1 + \alpha^3$$
$$\left(3 - 3\alpha^3\right)\alpha^2 = 1 + \alpha^3$$

slightly above .5 the C, C equilibrium for the group abruptly vanishes. To understand our proposed alternative let us step back for a moment to consider mixed strategy equilibria in ordinary finite games. There also the best response changes abruptly as beliefs pass through the critical point of indifference, albeit with the key difference that at the critical point randomization is allowed. But the abrupt change in the best response function still does not make sense from an economic point of view. A standard perspective on this is that of Harsanyi (1973) purification, or more concretely the limit of McKelvey and Palfrey (1995)'s Quantal Response Equilibria. Here the underlying model is perturbed in such a way that as indifference is approached players begin to randomize and the probability with which each action is taken is a smooth function of beliefs. In the limit as the perturbation becomes small, like the Cheshire cat, only the randomization remains. Similarly, in the context of group behavior, it makes sense that as the beliefs of a group change the probability with which they play different equilibria varies continuously. Consider for example $\alpha^3 = 0.499$ versus $\alpha^3 = 0.501$. In a practical setting where nobody actually knows α^3 does it make sense to assert that in the former case player 1 and 2 with probability 1 agree that $\alpha^3 \leq 0.5$ and in the latter case that $\alpha^3 > 0.5$? We think it makes more sense that they might agree that $\alpha^3 \leq 0.5$ with 90% probability and mistakenly agree that $\alpha^3 > 0.5$ with 10% probability in the first case and conversely in the second case. Consequently when $\alpha^3 = 0.499$ there would never-the-less be a 10% chance that the group would choose to play D, D not realizing that C, C is incentive compatible, while when $\alpha^3 = 0.501$ there would be a 10% chance that they would choose to play C, C incorrectly thinking that it is incentive compatible. We will develop below a formal model in which groups have beliefs that are a random function of the true play of the other groups and are only approximately correct. For the moment we expect, as in Harsanyi (1973), that in that limit only the randomization will remain. Our first step is to introduce a model that captures the grin of the Cheshire cat - we will simply assume that randomization is possible at the critical point. In the example we assert that when $\alpha^3 = 0.5$ and the incentive constraint exactly binds, the equilibrium "assigns" a probability to C, C being the equilibrium that is chosen by the group.[5] That is, when the incentive constraint holds exactly we do not assume that the group can choose their most preferred equilibrium, but instead we assume that there is an endogenously determined probability that they will be able to choose that equilibrium.

Remark. Discontinuity and non-existence is not an artifact of restricting attention to Nash equilibrium. The same issue arises if we assume that players 1 and 2 can use correlated strategies. When the game is a PD, that is, $\alpha^3 > 1/2$ then strict dominance implies that the unique Nash equilibrium is also the unique correlated equilibrium. When $\alpha^3 \leq 1/2$ the Nash equilibrium at C, C Pareto dominates every other correlated strategy, hence remains

[5]This is similar to Simon and Zame (1990)'s endogenous choice of sharing rules.

the unique best choice for players 1 and 2. When $\alpha^3 \leq 1/2$ the correlated equilibrium set is indeed larger than the Nash equilibrium set (containing at the very least the public randomizations over the Nash equilibria), but these correlated equilibria are all inferior for players 1 and 2 to C, C so will never be chosen. While it is true that the correlated equilibrium correspondence is better behaved than the Nash equilibrium correspondence - it is convex valued and upper-hemicontinuous - this example shows that the selection from that correspondence that chooses the best equilibrium for the group is never-the-less badly behaved - it is discontinuous. It is well known from the earliest work on competitive equilibrium, Arrow and Debreu (1954) that for the best choice from a constraint set to be well-behaved the constraint set needs to be lower-hemicontinuous and neither the Nash nor correlated equilibrium correspondence satisfies that property.

Remark. This example also explains our use of the term "game among groups". The game is actually a game among individual players, with actions and payoffs specified accordingly. But the subsets of players we call groups can act collusively by choosing among group profiles which are equilibria within the group. To illustrate, in the above game the only Nash equilibrium is D, D, D. No player can profitably deviate from that profile individually - but the group of 1 and 2 would deviate to C, C, which is better and incentive compatible within the group, if 3 played D.

3. Exogenous Groups and Collusion Constrained Equilibrium

We now introduce our model of exogenously specified homogeneous groups in which the groups pursue their own interest subject to within-group individual incentive constraints.

There are players $i = 1, 2, \ldots I$ and groups $k = \mathbf{1}, \mathbf{2}, \ldots \mathbf{K}$. The actions available to a player depend entirely on which group he is in; actions available for members of group k are A^k, assumed to be a finite set. We assume that there is a fixed assignment of players to groups $k(i)$. Notice that each individual is assigned to exactly one group and that the assignment is exogenous. All players within a group are symmetric - that is the groups are homogeneous - so the relevant utility of player i is $u^{k(i)}(a^i, a^{-i})$ and is invariant with respect to within group permutations of the labels of other players within their respective groups. If we let $\mathcal{A}^k$ denote the mixed actions for a member of group k, profiles of play chosen from this set represent the universe in which in-group equilibria reside. As should be clear from the motivating example, we will need to consider randomizations over in-group symmetric equilibria: each group is assumed to possess a private randomizing device observed only by members of that group that can be used to coordinate group play.

Because $\mathcal{A}^k$ is infinite, randomization over this set by the group leads to technical and conceptual complications that we prefer to avoid, so we will restrict the set of possible choices for the group. Specifically, we fix a finite subset $A^{kR} \subseteq \mathcal{A}^k$ containing all pure strategies and possibly mixed actions as well, and consider only in-group equilibria for group k in which all players choose the same action $a^k \in A^{kR}$. For example, with $A^k = \{H, T\}$ the actions in A^{kR} can be of the form: choose H, choose T, or randomize 50-50 between H and T. In other words, the model is consistent with independent individual randomization provided that individuals are limited to a finite grid of probabilities. Since in-group mixed equilibria may not be present in A^{kR} we will allow the group to choose in-group ϵ-equilibria in which small violations of the incentive constraints are allowed.[6]

Given the symmetry restriction we can simplify notation and write $u^k(a^i, a^k, \alpha^{-k})$ for the expected utility of player i in group $k(i) = k$ when a^i is his choice, the other group members play the common group action $a^k \in A^{kR}$, and any other group $\kappa \neq k$ assigns probability $\alpha^\kappa(a^\kappa)$ to all members of the group playing $a^\kappa \in A^{\kappa R}$.

Further, since only deviations from the common strategy matter, for player i in group $k(i) = k$ we need not allow a^i to take values in all of A^{kR} - it is sufficient to consider $a^i \in A^k \cup \{a_0^k\}$ where a_0^k means: "play the common mixed action $a^k \in A^{kR}$". That is, it is enough to consider deviations by player i to pure strategies A^k, letting $u^k(a_0^k, a^k, \alpha^{-k}) = u^k(a^k, a^k, \alpha^{-k})$ to be the utility when no deviation has taken place. Not only does this potentially greatly reduce the set of a^i that need be considered, but extends in a straightforward way when we come to consider correlated group strategies below. Notice that this formulation incorporates the use of randomizing devices that are private to the group: member i knows the result of the own group randomization $a^{k(i)}$ when choosing a^i, but does not know results of the randomization by other groups. This allows a limited amount of correlation within the group: they have a private randomizing device. We discuss the possibility that they engage in more elaborate correlation in Appendix 1.

Groups are assumed to be collusive - but they may collude only to choose plans that respect individual incentive constraints. The key reason that we start by considering homogeneous groups is that since group members are *ex ante* identical there is an "obvious" group objective, which is to assume that all members are treated equally and that the objective of the group is to maximize the common utility that they receive when all play the same action.

As indicated we allow a small amount of slack in the individual incentive constraints. Specifically, we introduce strictly positive numbers $v^k > 0$ measuring in utility units the violation of incentive constraints that are allowed. For a mixed profile α^{-k} by other groups and an action $a^k \in A^{kR}$ by group k we may define the gain function $G^k(a^k, \alpha^{-k}) = \max_{a^i \in A^k \cup \{a^k\}}[u^k(a^i, a^k, \alpha^{-k}) - u^k(a^k, a^k, \alpha^{-k})]$ as the degree to which the incentive constraint is violated at a^k (the smaller the gain the more stable the action). When the gain is

[6] This universe might encompass correlated strategies for the group, as indeed is the case in the introductory example, where the group mixes 50-50 between C, C and D, D. For expositional simplicity we formally deal with correlation in Appendix 1. Indeed, Allowing groups to use correlated equilibria within groups does not change the subsequent results, it merely requires a different and slightly more complex notation.

strictly less than v^k then a^k must be chosen by the group if it is to the benefit of the group to do so. When the gain is greater than v^k then a^k the group cannot choose a^k. When the gain is exactly v^k then the group may mix with any probability onto a^k. This is the same Cheshire grin logic as in the example, except that in the example we took $v^k = 0$.

Define $U^k(\alpha^{-k}) = \max_{\{a^k \in A^{kR} | G^k(a^k, \alpha^{-k}) < v^k\}} u^k(a_0^k, a^k, \alpha^{-k})$ to be the most utility attainable against α^{-k} when the incentive constraints are violated by strictly less than v^k (equal to $-\infty$ if the constraint set is empty). Then we take the finite set

$$B^k(\alpha^{-k}) = \{a^k \in A^{kR} | G^k(a^k, \alpha^{-k}) \leq v^k, u^k(a_0^k, a^k, \alpha^{-k}) \geq U^k(\alpha^{-k})\}$$

to represent actions that are feasible for the group given α^{-k}. We refer to this as the *shadow response set.* They are actions which violate the incentive constraints by strictly less than v^k and yield $U^k(\alpha^{-k})$, the most possible among such actions, plus those actions with $G^k(a^k, \alpha^{-k}) = v^k$ that yield at least $U^k(\alpha^{-k})$ - but possibly more. Observe that not all actions in $B^k(\alpha^{-k})$ need be indifferent, but that on the other hand all incentive compatible actions outside of $B^k(\alpha^{-k})$ are strictly worse for the group than any of those inside $B^k(\alpha^{-k})$.

Definition 1. A *collusion constrained equilibrium* is an α^k for each group that places weight only on $B^k(\alpha^{-k})$.

Define $\overline{B}^k(\alpha^{-k}) = \arg\max_{\{a^k | G^k(a^k, \alpha^{-k}) \leq v^k\}} u^k(a^k, a^k, \alpha^{-k}) \subseteq B^k(\alpha^{-k})$ to be the set of actions that maximize utility subject to the incentive constraints. Again, the key to collusion constrained equilibrium is that we allow a positive probability of actions in $B^k(\alpha^{-k})$ not merely in $\overline{B}^k(\alpha^{-k})$. If in a collusion constrained equilibrium α^k places positive weight on $B^k(\alpha^{-k}) \backslash \overline{B}^k(\alpha^{-k})$ we say that group k engages in *shadow mixing,* meaning that it is putting positive probability on alternatives it is not indifferent to. This may occur when best alternatives are not *strictly* incentive compatible, hence - this is our rationale for this equilibrium - they are not available to play with certainty within the group. This is to be contrasted with putting weight on $\overline{B}^k(\alpha^{-k})$ which are mixtures in the normal sense of indifference. Our example above shows that shadow mixing may be necessary in equilibrium.

Example 2. To illustrate the definition we apply it to the game of Example 1. If player 3 plays C with probability α^3 and the group plays D, D a player in the group who deviates to C gets $\alpha^3(-2) + (1 - \alpha^3)(-1)$ so this deviation is never profitable, D, D being strictly incentive compatible. If the group plays C, C the player who deviated to D gets $\alpha^3 \cdot 2 + (1 - \alpha^3) \cdot (-2) = 2(2\alpha^3 - 1)$: the best in-group equilibrium is thus incentive compatible for $2(2\alpha^3 - 1) \leq v^1$, at equality incentive compatibility is just satisfied and the equilibrium vanishes for larger values. So the condition for shadow mixing between C, C and D, D is $2\alpha^3 - 1 = v^1/2$ or $\alpha^3 = (1 + (v^1/2))/2$. Formally, for this value of α^3 the shadow response set $B^1(\alpha^3) = \{C, D\}$ for D is the only, hence best, action satisfying incentive compatibility strictly. For player 3 to be indifferent between C and D, letting p be the probability with which the group plays C, C we get the condition $5p = 5(1 - p)$ so $p = 1/2$. So equilibrium is that the group mixes

50-50 between C, C and D, D and player 3 plays C with probability $\alpha^3 = (1 + v^1/2)/2$. As $v^1 \to 0$ this converges to 1/2 which is the equilibrium in the original example.

The assumption that $v^k > 0$ plays a dual role in the model. First as indicated, we need to allow positive v^k if we wish to insure that in-group mixed equilibria are not excluded.[7] However, $v^k > 0$ plays a second role: it enables us to properly allow mixing only at "critical" points where small changes in beliefs lead to a discontinuous change in behavior.

Example 3. Group **1** has three actions H, M, L while group **2** has two actions H, L. For player i in group $k(i) = \mathbf{2}$ payoffs are $u^{\mathbf{2}}(a^i, a^{\mathbf{2}}, a^{\mathbf{1}}) = 0$, so group **2** has no active role and we concentrate on group **1**. For player i in group $k(i) = \mathbf{1}$ payoffs $u^{\mathbf{1}}(a^i, a^{\mathbf{1}}, a^{\mathbf{2}})$ are in the following matrix:

	$a^i = H, a^{\mathbf{2}} = H, L$	$a^i = M, a^{\mathbf{2}} = H, L$	$a^i = L, a^{\mathbf{2}} = H$	$a^i = L, a^{\mathbf{2}} = L$
$a^{\mathbf{1}} = H$	2	2	3	1
$a^{\mathbf{1}} = M$	1	0	1	1
$a^{\mathbf{1}} = L$	1	1	1	1

Action M is never part of an equilibrium: whatever the other group is doing, if the other members of your group play M you want to deviate. On the other hand no one ever wants to deviate from L - but incentive constraints are satisfied with exact equality there. Behavior against H is richer: you may want to deviate if the other group is playing H with high enough probability . Specifically, equilibria are computed to be as follows. Let $\alpha^{\mathbf{2}}$ be the probability with which members of group **2** play H, and observe first that any $\alpha^{\mathbf{2}}$ is an equilibrium for group **2**.

If $\alpha^{\mathbf{2}} \leq 1/2$ there are two in-group equilibria for group **1**: H and L; if $\alpha^{\mathbf{2}} > 1/2$ the only in-group equilibrium is L. In all in-group equilibria the incentive constraints are exactly satisfied (when $\alpha^{\mathbf{2}} \leq 1/2$ and group **1** action is H action M gives you the same utility as H; this is the role of M in the example).

So given the mixing rule we have specified above, with $v^{\mathbf{1}} = 0$ the collusion constrained equilibria consist of $\alpha^{\mathbf{2}} \leq 1/2$ and any vector $\alpha^{\mathbf{1}} = (a, 0, b)$, and $\alpha^{\mathbf{2}} > 1/2$ together with $\alpha^{\mathbf{1}} = (0, 0, 1)$. The group cannot guarantee that it will collude on the preferred action H.

With $v^{\mathbf{1}} > 0$ observe that $2 + v^{\mathbf{1}} = (1/2 + v^{\mathbf{1}}/2) \cdot 3 + (1/2 - v^{\mathbf{1}}/2) \cdot 1$ so that members of group **1** are indifferent between the payoff $2 + v^{\mathbf{1}}$ they get from agreeing with the group at H and deviating to L against group **2** playing $\alpha^{\mathbf{2}} = 1/2 + v^{\mathbf{1}}/2$.[8] Hence the collusion constrained equilibria consist of: (1) $\alpha^{\mathbf{2}} < 1/2 + v^{\mathbf{1}}/2$ and $\alpha^{\mathbf{1}} = (1, 0, 0)$, where H is strictly incentive compatible and best group alternative; (2) $\alpha^{\mathbf{2}} = 1/2 + v^{\mathbf{1}}/2$ and any vector $(a, 0, b)$, where the only strictly incentive compatible action is L hence $B^{\mathbf{1}}(\alpha^{\mathbf{2}}) = \{H, L\}$; and (3) $\alpha^{\mathbf{2}} > 1/2 + v^{\mathbf{1}}/2, \alpha^{\mathbf{1}} = (0, 0, 1)$. As we see, for $\alpha^{\mathbf{2}}$ slightly larger than 1/2 incentive constraints are violated but the violation is small enough to make collusion on H viable.

[7] The importance of this issue is underscored by the possibility of a unique in-group equilibrium, which is mixed.

[8] We are abusing terminology a bit: they do not "get $2 + v^{\mathbf{1}}$", but as long as the left member is larger than the right one the gain to deviating to L is less than $v^{\mathbf{1}}$.

Using $v^{\mathbf{1}} > 0$ captures the difference between $\alpha^{\mathbf{2}} < 1/2$ and the critical economy where a small change in $\alpha^{\mathbf{2}}$ makes H no longer viable. In a sense it captures the fact that indifference for $\alpha^{\mathbf{2}} < 1/2$ is not fundamental - it occurs just because there is an action M to which individuals are indifferent - but small perturbations in $\alpha^{\mathbf{2}}$ leave that indifference unchanged. Put differently, if we think that the inability of the group to coordinate perfectly is due to the fact that a small randomization in beliefs about the other group may cause indifference to be violated, then the "razor edge" equilibria for $\alpha^{\mathbf{2}} < 1/2$ are not vulnerable while the critical economy at $1/2 + v^{\mathbf{1}}/2$ is and this is correctly picked up when we make $v^{\mathbf{1}}$ strictly positive.

Example 4. This example is in Web Appendix 1 to save space. In this example the notion of collusion constrained equilibrium captures how "keeping up with the Jones's" type preferences may be a problem not at the individual level but at the level of the family. The example highlights two things: (1) the (Nash) equilibrium selection role that (pure) collusion constrained equilibrium can play; and (2) how the inability to commit to not collude may lead to inferior outcomes.

Incentive Compatible Games

There are two kinds of mixing: the group can mix between different actions chosen by the group using the group randomization device, but also individuals can mix. As we noted above individual mixing is included in the finite set A^{kR}, so the group mixes over a finite rather than continuous set. From an economic and empirical point of view dealing with approximate equilibria within the group does not pose a problem - in the field, laboratory or computationally we cannot expect individuals to achieve more than an approximate equilibrium.

If A^{kR} contains a relatively fine grid of mixtures there will be an ϵ-Nash equilibrium with a small value of ϵ. As long as v^k is strictly bigger than ϵ the group can find an action that is guaranteed to satisfy the incentive constraints to the required degree. Specifically, define $g^k = \max_{\alpha^{-k}} \min_{a^k \in A^{kR}} G^k(a^k, \alpha^{-k})$ so that regardless of the behavior of the other groups there is always a g^k approximate equilibrium within the group.

Definition 2. A game is *incentive compatible* if $v^k > g^k$ for all k.

Hereafter we will restrict attention to incentive compatible games: roughly this means that we chose a "fine enough" grid for each group.

4. Analysis of the Model

Having defined collusion constrained equilibrium we now want to show that they exist and make sense. In this section we consider how collusion constrained equilibria arise as the limits of fully collusive equilibria with random group beliefs and analyze more closely the role of shadow mixing. In the next section we will consider a concrete non-cooperative game involving representative or virtual players from each group and show that it gives rise exactly to collusion constrained equilibria.

4.1. The Existence of Collusion Constrained Equilibria

In this subsection we show that the basic problem of non-existence that arises when group try to choose actions in $\overline{B}(\alpha^{-k})$ is resolved by collusion constrained equilibrium.

Theorem 1. *In an incentive compatible game a collusion constrained equilibrium exists.*

This result follows from the following basic property of the shadow response set:

Lemma 1. *(i) In an incentive compatible game $\overline{B}^k(\alpha^{-k})$ is non-empty for all α^{-k}; (ii) every α^{-k} has an open neighborhood $\mathcal{A}$ such that $\tilde{\alpha}^{-k} \in \mathcal{A}$ implies that $B^k(\tilde{\alpha}^{-k}) \subseteq B^k(\alpha^{-k})$.*

Proof. Assertion (i) is obvious from the definition. (ii) If not there must be a sequence $\alpha_n^{-k} \to \alpha^{-k}$ and points $a_n^k \in B(\alpha_n^{-k}), a_n^k \notin B(\alpha^{-k})$. Since A^{kR} is a finite set, we may assume that we have chosen a subsequence along which $a_n^k = a^k$ is constant. Since G^k is continuous in α_n^{-k} any a^j such that $G^k(a^j, \alpha^{-k}) < v^k$ satisfies $G^k(a^j, \alpha_n^{-k}) < v^k$ for n large enough, so since A^{kR} is finite all those which satisfy the constraint strictly in the limit do so for n large enough, which implies that for such n it is $U(\alpha_n^{-k}) \geq U(\alpha^{-k})$. Let $\tilde{a}^k \in \arg\max_{\{a^k | G^k(a^k, \alpha^{-k}) < v^k\}} u^k(a^k, a^k, \alpha^{-k})$. Then $U^k(\alpha^{-k}) = u^k(\tilde{a}^k, \tilde{a}^k, \alpha^{-k})$ and since $a^k \in B(\alpha_n^{-k})$ for all n we then have

$$u^k(a^k, a^k, \alpha^{-k}) \geq u^k(\tilde{a}^k, \tilde{a}^k, \alpha^{-k}) = U^k(\alpha^{-k}).$$

By continuity of G^k it is also the case that $G^k(a^k, \alpha^{-k}) \leq v^k$ so we obtain $a^k \in B(\alpha^{-k})$, a contradiction. □

Proof of Theorem 1. Call $C(\alpha^{-k})$ the set of distributions over $B(\alpha^{-k})$. A profile α is a collusion constrained equilibrium if $\alpha^k \in C(\alpha^{-k})$ for all k, that is if $\alpha \in C(\alpha) \equiv \times_k C(\alpha^{-k})$, in other words if α is a fixed point of the correspondence $\alpha \twoheadrightarrow C(\alpha)$. Since the game is incentive compatible $C(\alpha^{-k})$ is non empty for any α^{-k}. Further, by construction, it is a convex valued correspondence. As a result, the correspondence $C(\alpha)$ is non empty and convex valued. By Lemma 1 we know that that $B(\alpha^{-k})$ is upper hemicontinuous. In turn this implies that both $C(\alpha^{-k})$ and $C(\alpha)$ are upper hemicontinuous. Hence the fixed point sought for exists by the Kakutani fixed point theorem. □

4.2. Random Beliefs

We now show that collusion constrained equilibria are limit points of standard equilibria when beliefs of each group about behavior of the other groups are random and randomness tends to vanish. We start by describing a random belief model. The idea is that given the true play α^{-k} of the other groups, there is a common belief $\tilde{\alpha}^{-k}$ by group k that is a random function of that true play. Notice that these random beliefs are shared by the entire group - we could also consider individual belief perturbations, but it is the common component that is of interest to us, because it is this that coordinates group play. Conceptually if we think that a group colludes through some sort of discussions that give rise to common knowledge - looking each other in the eye, a handshake and suchlike - then it makes sense that during these discussions a consensus emerges not just on what action to take, but underlying that choice, a consensus on what the other groups are thought to be doing. We must emphasize: our model is a model of the consequences of groups successfully colluding - we do not attempt to model the underlying processes of communication, negotiation and consensus that leads to their successful collusion.

Definition 3. An *ϵ-random group belief model* is a density function $f^k(\tilde{\alpha}^{-k}|\alpha^{-k})$ that is continuous as a function of $\tilde{\alpha}^{-k}, \alpha^{-k}$ and satisfies $\int_{|\tilde{\alpha}^{-k}-\alpha^{-k}|\leq\epsilon} f^k_\epsilon(\tilde{\alpha}^{-k}|\alpha^{-k})d\tilde{\alpha}^{-k} \geq 1-\epsilon$.

It is important to know that there are ϵ-random belief models for every positive value of ϵ. An obvious idea is to take a smooth family of probability distributions with mean equal to the truth and small variance. A good candidate for a smooth family is the Dirichlet since we can easily control the precision by increasing the "number of observations." However using an unbiased probability distribution will not work - it is ill-behaved on the boundary: if we try to keep the mean equal to the truth, then as we approach the boundary the variance has to go to zero, and on the boundary there will be a spike. A simple alternative is to bias the mean slightly towards a fixed strictly positive probability vector alpha with a small weight on that vector, and then let that weight go to zero as we take the overall variance to zero. The next example shows that this works.

Example 5. Let M^{-k} be the number of actions in A^{-k} and set $h(\epsilon) = (\epsilon/2)^2M^{-k}/(M^{-k} - (\epsilon/2)^2)$. Fix a strictly positive probability vector over A^{-k} denoted by β^{-k} and call the *ϵ-Dirichlet belief model* the Dirichlet distribution with parameters

$$\frac{1}{h(\epsilon)}\left[(1-\frac{\epsilon}{2\sqrt{2}})\alpha^{-k}(a^{-k}) + \frac{\epsilon}{2\sqrt{2}}\beta^{-k}(a^{-k})\right]$$

Theorem 2. *The ϵ-Dirichlet belief model is an ϵ-random belief model.*

Proof. Since the parameters are away from the boundary by at least $\epsilon/2$ this has the requisite continuity property. It has mean $\overline{\alpha}^{-k} = (1-\frac{\epsilon}{2\sqrt{2}})\alpha^{-k} + \frac{\epsilon}{2\sqrt{2}}\beta^{-k}$. Set $\hat{\alpha}^{-k} = (1-\frac{\epsilon}{2\sqrt{2}})\tilde{\alpha}^{-k} + \frac{\epsilon}{2\sqrt{2}}\beta^{-k}$. Since the covariances of the Dirichlet are negative, $E|\hat{\alpha}^{-k} - \overline{\alpha}^{-k}|^2$ is bounded by the sum of the variances and we may apply Chebyshev's inequality to find

$$Pr[|\hat{\alpha}^{-k} - \overline{\alpha}^{-k}| > \epsilon/2] \leq E|\tilde{\alpha}^{-k} - \beta^{-k}|^2/(\epsilon/2)^2 \leq M^{-k}h(\epsilon)/[\epsilon(M^{-k}+h(\epsilon))] \leq \epsilon/2.$$

Observe that $|\hat{\alpha}^{-k} - \overline{\alpha}^{-k}| = (1-\frac{\epsilon}{2\sqrt{2}})|\tilde{\alpha}^{-k} - \beta^{-k}| \geq |\tilde{\alpha}^{-k} - \beta^{-k}| - \frac{\epsilon}{2}$. Hence $Pr(|\tilde{\alpha}^{-k} - \beta^{-k}| > \epsilon) \leq \epsilon/2 \leq \epsilon$ which shows that this is indeed an ϵ-random belief model. □

Fix some probability distribution $F^k(\alpha^{-k})$ over $\overline{B}^k(\alpha^{-k})$ measurable as a function of α^{-k}. Define $R^k(a^k|\alpha^{-k}) = \int F^k(\tilde{\alpha}^{-k})[a^k]f^k(\tilde{\alpha}^{-k}|\alpha^{-k})d\tilde{\alpha}^{-k}$. Notice that for given beliefs $\tilde{\alpha}^k$ we are assuming that the group colludes on a response in $\overline{B}^k(\tilde{\alpha}^{-k})$ which are the best choices for the group that weakly satisfy the incentive constraints, and not on points in

$B^k(\tilde{\alpha}^{-k})\backslash\overline{B}^k(\tilde{\alpha}^{-k})$ as would be permitted by shadow mixing. We define an *ϵ-random belief equilibrium* as an α_ϵ such that $\alpha_\epsilon^k = R^k(\alpha_\epsilon^{-k})$. The key result is

Theorem 3. *Fix a family of ϵ-random group belief models, an $F^k(\alpha^{-k})$ and an incentive compatible game. Then for all $\epsilon > 0$ there exist ϵ-random belief equilibria. Further, if α_ϵ are ϵ-random group equilibria and $\lim_{\epsilon\to 0}\alpha_\epsilon = \alpha$ then α is a collusion constrained equilibrium.*

Proof. By the Lebesgue dominated convergence theorem R^k is continuous, so we may apply the Brouwer fixed point to get existence of ϵ-random group equilibria. Now consider a sequence of ϵ-random group equilibria with $\lim_{\epsilon\to 0}\alpha_\epsilon = \alpha$. By Lemma 1 we know that for sufficiently small ϵ, $|\alpha_\epsilon^{-k} - \alpha^{-k}| \leq \epsilon$ implies $B^k(\alpha_\epsilon^{-k}) \subseteq B^k(\alpha^{-k})$. Hence for such α_ϵ^k and ϵ it must be that $\alpha_\epsilon^k(B^k(\alpha^{-k})) = 1$ with $\alpha^k(B^k(\alpha^{-k})) = 1$ at the limit - which is the condition for a collusion constrained equilibrium. □

We should emphasize that this result is not an equivalence result: random belief equilibria converge as $\epsilon \to 0$ to collusion constrained equilibria. However, there is no assertion that all collusion constrained equilibria arise this way. This is similar to the result for Harsanyi (1973) where convergence of random utility equilibria to Nash equilibria is assured, but only under additional conditions do we know that Nash equilibria arise as limits of random utility equilibria. In cases such as quantal response indeed, limits of quantal response equilibria are a refinement of Nash equilibrium.

4.3. When Does Shadow Mixing Matter?

For applications it is useful to know when groups do not engage in shadow mixing. There are two important cases where groups will engage only in ordinary mixing.

1. The action that maximizes group utility without constraint is always an in-group equilibrium. Since the action is an equilibrium, it strictly satisfies the relaxed constraint with $v^k > 0$. Since it maximizes group utility without any constraint, it certainly maximizes group utility with the constraint, so $\overline{B}^k(\alpha^{-k}) = B^k(\alpha^k)$. Notice that in case the group has a single player, or more generally the game is a game of common interest so that group members always get the same payoffs as each other regardless of the actions chosen this assumption is satisfied.[9]

2. Separable games in which $u(a^i, a^k, a^{-k}) = w(a^{-k}) - c(a^i, a^k)$ so that the incentive constraints do not depend on what the other groups do. Here $G(a^k, \alpha^{-k}) = \max_{a^i \in A^k} c(a^k, a^k) - c(a^i, a^k)$ independent of α^{-k}. Hence for generic v^k there will be no a^k for which $G(a^k, \alpha^{-k}) = v^k$. These models can be important for applications because they can be thought of as approximations in political economy games such as voting or lobbying games where the group

[9]In these games an action profile maximizing the utility of some group member does the same for each group member and must therefore be an in-group equilibrium too.

size is large so individuals perceive that their own action has no impact on the common public good w - for example, the outcome of a vote.

4.4. What Difference Do Collusion Constraints Make?

We return to example 1 to illustrate how accounting for incentive and collusion constraints may impact on the strategic analysis of a game.

First, the only Nash equilibrium of the game consists of all players to play D. To see this observe that as shown in Footnote 4 players 1 and 2 can mix only if $\alpha^3 \leq 1/2$ and then $\alpha^1 = \alpha^2$ are increasing in α^3; so the smallest value of α^1 occurs when $\alpha^3 = 0$ and it is $\alpha^1 = 1/3$. But for $\alpha^1 = \alpha^2 \geq 1/3$ player 3's best response is to play C for sure; hence there is no equilibrium in which player 1 and 2 mix. The two of them playing C, C is not an equilibrium because 3's best response to it is C for sure, but in that case they will play D, D. Profile D, D, D on the other hand is Nash. In this equilibrium payoffs are $(5, 5, 5)$.

On the other hand, ignoring individual incentive constraints, that is assuming that the group will collude on best group action, leads to predict that players 1 and 2 will play C, C in which case 3 also chooses C. Predicted payoffs would be $(6, 6, 5)$.

Consider now collusion constrained equilibrium. We have seen in Example 2 that in this equilibrium the group mixes 50-50 between C, C and D, D and player 3 plays C with probability $\alpha^3 = (1 + \epsilon/2)/2$. In equilibrium player 3 gets 2.5. Players 1 and 2 get $4(\frac{1}{2} + \frac{\epsilon}{4}) + \frac{11}{2}(\frac{1}{2} - \frac{\epsilon}{4}) = 4\frac{3}{4} - \frac{3}{8}\epsilon$. As $\epsilon \to 0$ the limit payoff vector is a much lower$(4.75, 4.75, 2.5)$.

As can be expected, ignoring individual constraints lead to an unrealistically optimistic conclusion. But the remarkable point is that in the example the same is true for Nash equilibrium: ignoring collusion constraints also leads to predicting higher utilities for the players. Incidentally, this is why we call our equilibrium collusion *constrained*: in general collusion makes the group of the whole worse off.

Notice that a benevolent mechanism designer who could choose between having players play the game and a safe alternative that gave payoffs of $(4.9, 4.9, 4.9)$ who either analyzed the game ignoring collusion or who analyzed the game assuming that players could collude would choose the game over the safe alternative, while a designer who recognized that collusion is subject to incentive constraints would reach the opposite conclusion.

Given that the collusion constrained payoff is smaller than Nash payoff, a natural question is, why do group players not wave goodbye to the group and play Nash? The problem is that this would not be credible. The groups we have in mind cannot be prevented from colluding, that is, their members cannot credibly commit not to collude. For example how can farmers stop interacting with other farmers to credibly commit to not acting as a farm lobby? In our case, whatever players 1 and 2 declare about group membership, if player 3 plays D they will then want to collude on C, C and anticipating this player 3 will not play D in the first place.

5. Leadership Equilibrium

To give a concrete way in which collusion constrained equilibria can arise, we give a non-cooperative model of leadership which gives rise to collusion constrained equilibria. Leaders lead their group to act when several groups interact - they tell their followers things such as "let's go on strike" or "let's vote against that law." The idea is that group leaders serve as explicit coordinating devices for groups - and we will model them in a way that gives rise exactly to collusion constrained equilibrium. Each group will have a leader who tells group members what to do, and since he is to serve as an effective coordination device for group members these instructions cannot be optional for group members. However, we do not want leaders to issue instructions that members would not wish to follow - that is, that are not incentive compatible. Hence we give them incentives to issue instructions that are incentive compatible by allowing group members to "punish" their leader. As in the previous section incentive compatibility will mean that constraints can be violated by no more than v^k, and here this value has a concrete interpretation as the leader's "valence": the higher v^k the more members are ready to give up to follow the leader. Alternatively, v^k can be thought of as measuring group loyalty.

While this is intended as an abstract model of how groups can reach decisions, we observe that in fact it is often the case that groups follow orders given by a leader but engage in *ex post* evaluation of the leader's performance. Specifically, we will consider the following non-cooperative game. Each group is represented by two virtual players: a *leader* and an *evaluator,* each of whom has the same underlying preferences as the group members. Each leader has a punishment cost $P^k > \max_{a^j,a^k,a^{-k}} u^k(a^j, a^k, a^{-k}) - \min_{a^j,a^k,a^{-k}} u^k(a^j, a^k, a^{-k})$. The game goes as follows:

Stage 1: Each leader privately chooses an action plan $a^k \in A^{kR}$: conceptually these are orders given to the members who must obey the orders. All members of group k thus play a^k.

Stage 2: In each group, the evaluator observes the action plan of the leader and chooses a response a^i [10]

Payoffs: The evaluator receives utility $u^k(a^i, a^k, a^{-k}) + v^k \cdot I(a^i = a^k)$ where I is the indicator function, that is he gets the v^k bonus only if he chooses a^k. As to the leader, if the evaluator chooses a^i he gets $u^k(a^k, a^k, a^{-k}) - P^k I(a^i \neq a^k)$, that is, he receives the penalty P^k if and only if the evaluator disagrees with his decision. Note that the leader and evaluator do not learn what the other groups did until the game is over.

[10] The evaluation need not be done by a single evaluator, but by consensus or some other aggregation method by all or a subset of group members. It makes no difference to the results.

Theorem 4. *In an incentive compatible game* α *are sequential equilibrium choices by the leaders if and only if* $\alpha^k(a^k) > 0$ *implies* $a^k \in B^k(\alpha^{-k})$*, that is,* α *is a collusion constrained equilibrium.*

Proof. The key implication of sequentiality - see for example Fudenberg and Tirole (1991) - is that the beliefs of the evaluator about the mixtures of other leaders must be independent of the signal received from his own leader - since his leader has no information about the signals of the other leaders. Suppose first that α is sequential. Then the belief of the evaluator for group k about other groups is α^{-k}, independent of the signal received from his own leader - so in effect from the perspective of the evaluator this is treated as a constant.

Because the game is incentive compatible, the leader can insure himself a utility of $U^k(\alpha^{-k})$ by choosing the best a^k that strictly satisfies the incentive constraints since he will not be deposed in that case. If he makes an announcement that violates the incentive constraints he is deposed with probability one and gets $\underline{u}^k < U^k(\alpha^{-k})$, so it must be that any announcement with $\alpha^k(a^k) > 0$ has $a^k \in B^k(\alpha^{-k})$.

Suppose conversely that any announcement with $\alpha^k(a^k) > 0$ has $a^k \in B^k(\alpha^{-k})$. There are two kinds of $a^k \in B^k(\alpha^{-k})$: those for which the incentive constraints hold exactly and those for which they hold strictly. If they hold strictly, then the benevolent leader gets $U^k(\alpha^{-k})$ by the definition of U^k. If they hold weakly, then the evaluator is indifferent between choosing a^k and keeping the leader and picking an alternate best response and penalizing him. Hence the probability that the leader is penalized $p^k(a^k, \alpha^{-k})$ may be any number between zero and one, and in particular may be chosen so that $u^k(a^k, a^k, \alpha^{-k}) - p^k(a^k, \alpha^{-k})P^k = U^k(\alpha^{-k})$ since by definition of B^k we have $u^k(a^k, a^k, \alpha^{-k}) \geq U^k(\alpha^{-k})$. This means the leader is indifferent between all actions in $B^k(\alpha^{-k})$ and in particular it is optimal for him to choose α^k since that places weight only on $B^k(\alpha^{-k})$. □

Remark. Provided that the penalty $P^k > \max_{a^j, a^k, a^{-k}} u^k(a^j, a^k, a^{-k}) - \min_{a^j, a^k, a^{-k}} u^k(a^j, a^k, a^{-k})$ the exact size of the penalty does not matter to the sequential equilibrium strategies of the leaders, α: this follows directly from Theorem 4 because the set of collusion constrained equilibria is defined without reference to P^k.

6. Choice of Leader and Endogenous Formation of Alliances

In a purely mechanical way the results on exogenous homogeneous groups extend to heterogeneous groups (and to correlated equilibrium), as we show in Appendix 1. But heterogeneity in particular raises issues of interpretation: taking as exogenous the "objective function of the group" makes it possible to prove the relevant theorems, but where does that objective function come from? Here we consider endogenizing both the objective function and members of the groups.

In the case of exogenous groups we have two models: an abstract model in which groups collude to choose the "best" equilibrium for the group, and a concrete game between leaders whose followers do as they are told and evaluate the performance of the leader *ex post.* In the exogenous group case we showed that the game between leaders is a concrete realization of the group collusion model by showing that the two models yield the same equilibrium behavior by the groups. The game between leaders, however, naturally suggests a richer setting for analysis. Several leaders may, for example, need to compete for groups. This leads us to study how groups might be endogenously aggregated by such candidate leaders. One can think, for example, of leaders[11] whose preferences are a weighted average of those of two or more different groups, who wish these groups to form a heterogeneous coalition and adopt a common course of action; and that these "common" leaders have to compete with "parochial" group leaders who inherit group preferences and wish to lead their own groups - and why not, possibly other groups too. We therefore allow, for our subsequent analysis, the possibility that each group is approached by different candidate leaders, each of whom, in turn, may possibly approach and lead several groups. The initial homogeneous groups may thus coalesce into larger heterogeneous aggregations.

To model a situation like this the first question then is: How do leaders compete for groups? What will a leader say to convince a group to follow him? The answer we adopt in this paper is that his message is of the following form: "Do what I say, and you will get utility U". That is, to win a coalition leaders will as before recommend actions, but will in addition make utility bids declaring what payoff the involved groups will attain. Groups will choose the leaders who offer them the highest "reasonable" utility level. They follow as before the action the chosen leader recommends, but punish him if they think his utility bid is too high compared to what they expect to actually get. Thus leaders who manage to form a coalition that follows them are punished under "discontent" conditions analogous to what we had before.[12] To illustrate the basic ideas we again start with a simple example.

Example 6. This is a simple game between two groups of at least three members each, which we call the *conformist's prisoner's dilemma*. The two groups are symmetric with each other, and players choose between two actions C, D. If all players in each group choose the group action the individual payoffs are given by the additively separable prisoner's dilemma game

	C	D
C	$1 , 1$	$-\gamma , 1+\gamma$
D	$1+\gamma , -\gamma$	$0 , 0$

Individual preferences reflect a desire for conformity: an individual player gets the payoff determined by the common action minus a fixed strictly positive penalty if he fails to choose the group action.[13] This means that any pure choice of action by the group is incentive

[11] From here on by "leader" we mean "candidate leader" - we most often omit the adjective for brevity.

[12] Observe that this is not a model of elections, where an overall winner sets rules all players must follow. Different groups or coalitions or groups will generally select different leaders and each will act according to the prescriptions of the leader they choose. Payoffs accrue in the game among groups from the profile of actions of the different groups and members.

[13] We impose that the group is composed of at least three members so that the "group of all players except you" is a majority against you.

compatible, and enables us to focus more clearly on the relation between the two groups.

In the model of two exogenous homogeneous groups the outcome is clear: each group has the dominant action of D and the outcome is that this is what both groups do and all players receive 0. But is there somehow a way out of this deadlock? Indeed: why should not somebody who can speak to both groups point out the clear benefit to all from forming a single group and have them coordinate on C under his leadership? Unfortunately, the common group is susceptible to a similar problem: why does not a member of, say, group 1 propose that by separating from the common group and playing D all members of group 1 would receive $1+\gamma$ instead of 1. Of course if both groups do this, we are back to 0 and joining the combined group seems attractive again.

Our proposal, as mentioned earlier, is to explicitly consider leaders that recommend actions as before and make utility bids in an effort to form coalitions. Group members will choose the best bid - but we require that bids be credible in the sense that the expected utility group members receive when they choose the best bid should in fact be at least the utility they were promised. Just as in the exogenous group model we imposed (and continue to impose) the requirement that the instructions of the leaders be acceptable to the members by having the members evaluate the instructions *ex post,* we now impose the requirement that bids be credible through *ex post* evaluation by the members.[14]

To explain what we propose to do let us assume that there are three leaders: two group leaders with preferences inherited from their respective groups, and a common leader who cares about the average utility of all members of both groups. The group leaders send offers only to their own group; the common leader sends offers to both groups. We will show later that in equilibrium the group leaders always recommend D while the common leader always says C. We simply assume this for now. The interest lies in utility bids. Again for simplicity in this example, let us suppose that they may only bid utility of either 0 or $2(1-\epsilon)$ where ϵ is a small positive number. Note that the high bid is closer to the C,C payoff of 1 than the low bid. Group members follow the leader who bids the highest utility, and in case of a tie they follow their own group leader (we will adopt this tie-break rule throughout).

First let us see if there can be a pure strategy equilibrium. If the group leaders both bid $2(1-\epsilon)$ then in fact everyone gets 0 and they are clearly seen to be liars and would be punished for sure. If they both bid 0 the common leader can bid $2(1-\epsilon)$. In this case everybody actually receives 1 and - given that the leader is constrained to bid $2(1-\epsilon)$ or 0 the claim of $2(1-\epsilon)$ is more accurate than the alternative bid of 0 so the common leader should be regarded as telling the truth. But: in this case a group leader can bid $2(1-\epsilon)$ and not be punished, for the bid will be accepted, the result will be that his group gets $1+\gamma$ so he also should be regarded as telling the truth. So there is no equilibrium in pure strategies.

[14]By *ex post* we mean after the group has accepted the leader's offer, as opposed to assessing the credibility of leaders' offers *ex ante* before selecting one. As will be clear from the following, offers are, however, judged before payoffs accrue. A word on this may be useful, because we are assuming for instance that if a leader proposes "War, Win for sure" and the group chooses the leader they go to war, but if they think there is a non-negligible (in what sense will again be clear in the sequel) probability of defeat they punish the leader, even if in the end the war is won.

Let us then look for a mixed equilibrium. Suppose the group leaders each bid $2(1-\epsilon)$ with probability p and that the common leader bids $2(1-\epsilon)$ with probability q. An individual who accepts an offer of $2(1-\epsilon)$ from the common leader gets $-p\gamma + (1-p)\cdot 1$, for with probability p the other group's leader wins by bidding $2(1-\epsilon)$ and the other group will play D, while with probability $1-p$ the common leader wins the other group and the outcome is C, C. For the common leader to be indifferent between the two bids given that he will be evaluated *ex post* the expected utility received by the groups should be $1-\epsilon$. Indeed only then are both the bids of $2(1-\epsilon)$ and 0 equally accurate and hence the extent to which he may be punished can be determined endogenously to make him indifferent between the two bids. So it must be that $-p\gamma + (1-p) = 1-\epsilon$ or $p = \epsilon/(1+\gamma)$. Now we examine the optimal choice by the group leaders. An individual who gets an offer of $2(1-\epsilon)$ from a group leader accepts that offer, plays D and gets $q(1-p)(1+\gamma)$, for they get $1+\gamma$ only if the other group plays C, which occurs in the event the common leader bids $2(1-\epsilon)$ and the other group leader bids 0. In order for the group leader to be willing to mix this utility must again be equal to $1-\epsilon$. That is, the condition for equilibrium is $q(1-p)(1+\gamma) = 1-\epsilon$. Substituting the equilibrium value of p we then get $q = (1-\epsilon)/(1+\gamma-\epsilon)$. Thus for small ϵ the equilibrium is approximately $q = 1/(1+\gamma)$ and $p = 0$.

This shadow mixing equilibrium seems to have sensible qualitative properties. The parameter γ measures how attractive defection is relative to cooperation. Cooperation occurs in equilibrium when common leader wins both groups, so equilibrium probability of cooperation is $q(1-p)^2$. When γ is small the conflict between the groups is small and the common group forms with high probability since $q(1-p)^2$ is near 1 and the groups cooperate most of the time. When γ is large the conflict between the groups is large, the common group forms with low probability since $q(1-p)^2$ is near 0 and the groups rarely cooperate.

6.1. A Model of Endogenous Alliances

The above analysis is limiting in a number of ways. For example: why should leaders be restricted only to make two offers - while it might make sense that they are limited to a finite set, it seems likely that they can make more refined offers than 0 or $2(1-\epsilon)$. Similarly why those two particular offers? What if the groups aren't conformist? What if group leaders can talk to both groups? And so forth. Here we introduce a more general model that captures the logic of the example while dropping the arbitrary limitations. After describing this model and basic results, we then use it to analyze a class of games which includes the conformist prisoner dilemma as a special case.

In the formal model leaders must induce groups to join them. They recommend actions and make utility bids. Members of a group choose the leader who makes the highest bid (accounting for valence), where this choice implies their commitment to follow the leader's recommendation for action. Credibility of his utility bid is then assessed by an evaluator and the leader is punished for lying as will be made precise shortly.

We continue with the framework that there are players $i = 1 \ldots I$ and groups $k = \mathbf{1}, \mathbf{2}, \ldots \mathbf{K}$ and that player i belongs to group $k(i)$, has available actions $A^{k(i)}$ and that the group has available actions in the finite set A^{kR}. We continue to write $u^{k(i)}(a^i, a^{k(i)}, a^{-k(i)})$ for the utility of a member of group $k(i)$. We continue to consider a collection of homogeneous groups. Now however, we consider a more flexible set of leaders.

There are leaders $\ell = 1, 2, \ldots, L$ where $L \geq K$. Each leader potentially leads an alliance: a set of groups $\mathcal{K}^\ell \subseteq \{1, \ldots K\}$ to which he can appeal. Leaders are assumed to have preferences of the form $u^\ell(a) = \sum_{k \in \mathcal{K}^\ell} \beta_k^\ell u^k(a)$ where β_k are some fixed non-negative weights with $\sum_{k \in \mathcal{K}^\ell} \beta_k = 1$. That is, a leader wishes to maximize some weighted average of the utility of the groups in his alliance. We do not attempt to explain where these weights come from, but we can consider for example competition between leaders of the same alliance who value the groups differently.

Each leader ℓ makes an offer $r^\ell = (a^{\ell k}, u^{\ell k})_{k \in \mathcal{K}^\ell}$ consisting, for each group k in his alliance, of an action to be played $a^{\ell k} \in A^{kR}$ and a utility level offered $u^{\ell k} \in U$. The utility offers are chosen from a common finite feasible set of bid utilities $U \subset \Re$, the same for all leaders. Let $\underline{u} = \min\{u | u \in U\}$. We assume that $\underline{u} < \min_{k, a^j, a^k, a^{-k}} u^k(a^j, a^k, a^{-k})$ and $\max\{u | u \in U\} > \max_{k, a^j, a^k, a^{-k}} u^k(a^j, a^k, a^{-k})$. The grid is evenly spaced with the gap between grid points equal to $d > 0$. We let R^ℓ be the set of all possible offers r^ℓ that can be made by leader ℓ.

As before leaders also have a valence $v^{\ell k}$ for $k \in \mathcal{K}^\ell$ which now may be different for different groups in his alliance. These valences are used both for evaluating actions and utility offers. We assume that no ties are possible for utility offers, that is, there is no group k, no pair of utilities $u, u' \in U$ and no pair of leaders $\ell \neq l$ with $k \in \mathcal{K}^\ell, \mathcal{K}^l$ such that $u + v^{\ell k} = u' + v^{lk}$.

We assume that each group gets offers from at least one leader. We assume that for each k there is a unique leader who has the highest valence $v^{\ell k}$ among those who are able to make offers to group k, and we denote this leader by $\ell = k$. We refer to the leader k as the *group leader* for k although that leader may also be able to make offers to other groups. In this context, if $v^{kk} > g^k$, that is if the group leaders have sufficient valences to find an incentive compatible plan for their group we say as in Section 3 that the game is *incentive compatible*. We consider incentive compatible games henceforth.

All leaders except for group leaders bidding to their own group can guarantee that they lose the bidding by bidding $\underline{u}$; and by assumption a losing bid is never punished. If ℓ (including group leaders) bids $u^k = \underline{u}$ to all groups $k \in \mathcal{K}^\ell$ we say that the leader has *opted-out*.

As in Section 5 each group has an evaluator. In addition to evaluating the action of the leader, the evaluator must now also evaluate the utility offer of the leader. He does so by choosing a *predicted utility* on the grid U. We assume that his payoff includes the square difference between predicted utility and actual utility. In the absence of a grid this would imply that his optimal choice is the conditional expectation of utility given his information. Given the constraint of the grid, if the conditional expectation is not exactly at the midpoint of a grid interval it is optimal for the evaluator to choose as predicted utility the unique point in U closest to the conditional expectation. In case expected utility is at the midpoint of a grid interval the evaluator is indifferent between choosing, as predicted utility, one of the two closest points on the grid and can mix between the two points. In either case the leader is punished if his bid is greater than the realized prediction of the evaluator, or if as in Section 5 the evaluator's chosen action is different than the one proposed by the leader. We assume that the punishments are cumulative: that is, if leader ℓ is punished by groups $k \in \mathcal{K}$ he suffers a penalty of $\sum_{k\in\mathcal{K}} P^{\ell k}$. We continue with the assumption that $P^{\ell k} > \max_{a^j,a^k,a^{-k}} u^k(a^j, a^k, a^{-k}) - \min_{a^j,a^k,a^{-k}} u^k(a^j, a^k, a^{-k})$.

The *alliance game* proceeds as follows:

Stage 1: each leader chooses an an offer $r^\ell \in R^\ell$

Stage 2: each group joins the alliance of the leader who offered that group the highest value of $u^{\ell k} + v^{\ell k}$ and the evaluator for group k observes the name ℓ of the alliance leader for his group and the entire offer r^ℓ of that leader

Stage 3: the evaluator for group k chooses a response (a^k, u^k), consisting of an action and a predicted utility

Payoffs: the evaluator gets $u^k(a^k, a^{\ell k}, a^{-k}) - (u^k - u^k(a^{\ell k}, a^{\ell k}, a^{-k}))^2 + I(a^k = a^{\ell k})v^{k\ell}$; if κ^ℓ are the groups that accepted leader ℓ's offer, then leader ℓ receives utility $u^\ell(a) - \sum_{k\in\kappa^\ell} I(a^k \neq a^{\ell k} \text{ or } u^k > u^{\ell k})P^{\ell k}$.

Note that if a group does not join a leader's alliance he still passively receives utility from them based on what they choose to do, according to the weight he gives to the group in his utility function. The evaluator does not observe the losing bids for his group: this is important because it means that - since the leader knows that the evaluator will evaluate his bid only if it is accepted - the evaluator evaluates the leader based on exactly the same information available to the leader. However, the evaluator does observe the bids made by the leader of the alliance he joins to other groups. This is also important: it ensures that if a leader deviates from the equilibrium path evaluators cannot speculate about what was offered by this leader to other groups.

6.2. Alliance Constrained Equilibrium and Correlated Equilibrium

We now define equilibrium for the above game and spell out a basic property of such an equilibrium.

Definition 4. A correlated strategy ρ over actions - that is a probability distribution derived from profiles of leaders recommendations and the rule for accepting highest bids - is an *alliance constrained equilibrium* if it corresponds to the strategies of the leaders in a sequential equilibrium in the alliance game.

By an ϵ-*correlated equilibrium* we mean a correlated strategy ρ over actions where the payoffs are those of the individual group members and nobody can benefit by more than ϵ by deviating.

Theorem 5. *In an incentive compatible game an alliance constrained equilibrium is a* $\max v^{k\ell}$*-correlated equilibrium.*

Proof. Each leader can guarantee that he is not punished. Non-group leaders can simply make minimum offers that will necessarily be rejected in favor of the group leaders, hence will not be punished. Group leaders can make minimal offers to all groups which means that none of those offers will be accepted except (possibly) the offer to their own group. For the offer to the own group the fact that the game is incentive compatible means that given the strategies of all other leaders there exists a recommendation for their own group that is strictly $v^{k\ell}$ incentive compatible resulting in no punishment.

Since every leader can guarantee that he is not punished, in equilibrium no offer can be punished with probability 1, meaning that conditional on the play of the other groups it must be at least $\max v^{k\ell}$ incentive compatible. □

Remark. We should highlight that our model of alliance constrained equilibrium involves a number of modeling choices. *a)* Evaluation is based on information available to the leader. If we were to assume otherwise we would get equilibria where leaders are forced to lie on the equilibrium path. *b)* We allow leaders to "lie" by understating utility. If we were to assume otherwise then leaders could be trapped into being punished for a high bid because they are unable to cut their bid slightly to avoid punishment. This would result in a plethora of equilibria and it is hard to see why they would make sense. *c)* We assume that punishments by different groups are cumulative, rather than, say assuming that the leader is punished if he is punished by some group who has a say over his punishment. This simplifies the analysis greatly because we do not need to take account of how the probabilities of punishments by different groups interact. *d)* The evaluators see the bids their leader made to other groups. If not they would have too much freedom to form implausible beliefs in response to deviations by leaders.

6.3. Equilibrium in the Prisoners Dilemma

We now study alliance constrained equilibria in the conformists prisoners dilemma of Example 6. We assume that the grid U starts at $\underline{u} < -\gamma$, has gaps of length $0 < d \le \gamma, 1/2$, and does not contain the points $d/2, 1+d/2$. We also let U_0 be the unique grid point in the open interval $(-d/2, d/2)$, and U_1 the unique grid point in $(1-d/2, 1+d/2)$.

The leadership structure is the one outlined in Example 6, we now describe it more precisely. There are three leaders, the group leaders $\ell = 1, 2$ with valences $v_1 = v_2 = v$ and a common leader $\ell = 3$ with valence $0 < v_3 < v$. The group leaders make offers only to their own groups and put all weight on their own group's utility. The common leader makes offers to both groups and puts equal weight on the utility of each group. Notice that for the common leader C strictly dominates D. We assume that $u - d + v < u + v_3$ so that if the group leaders underbid the common leader they lose; equivalently, $v - v_3 < d$.

Definition 5. A *strongly symmetric equilibrium* is one in which both group leaders play the same strategy, offer only D and bid greater than or equal to U_0 and in which the offers made by the common leader are of the form (C, u, C, u).

For a strongly symmetric equilibrium we denote by $R(u)$ the probability with which a group leader bids less than or equal to u and by $Q(u)$ the probability with which the common leader bids less than or equal to u. The following results are proved in Web Appendix 2.

Theorem 6. *In the PD game there is a unique strongly symmetric equilibrium. In this equilibrium no leader bids below U_0. For $U_0 \leq u < U_1$*

$$R(u) = \frac{u + \gamma + d/2}{1 + \gamma}$$

with $R(U_1) = 1$ and for $U_0 < u \leq U_1$

$$Q(u) = (\gamma/(\gamma + d))^{(U_1 - u)/d}$$

with

$$Q(U_0) = (\gamma/(\gamma + d))^{(U_1 - U_0 - d)/d} \left(\frac{\gamma}{U_0 + \gamma + d/2} \right).$$

Neither leader is punished for bidding U_0 and both leaders are punished with positive probability for each higher bid. In this equilibrium, each group cooperates with probability Π where

$$\frac{1 - d}{1 + \gamma} \leq \Pi \leq \frac{1}{1 + \gamma}.$$

It is worth noting that in equilibrium group leaders bid higher utility values with (weakly) decreasing probability, while the common leader bids higher values with increasing probability.

For comparative statics the continuum limit is cleaner to work with (proof in Web Appendix 2):

Theorem 7. *The limit of the unique strongly symmetric equilibrium as $d \to 0$ is given by*

$$R(u) = \frac{u + \gamma}{1 + \gamma}$$

with $R(1) = 1$, and for $0 \leq u \leq 1$

$$Q(u) = e^{(u-1)/\gamma}.$$

6.4. Transfer Games with Group and Common Leaders

A question naturally arises in the prisoners dilemma game: for general coalition constrained equilibria, can we bound the probability of defection from above? We answer this for a class of 2×2 *transfer games* which includes the conformists prisoners dilemma as a special case. The games in this class again have two conformist groups of at least three members each, symmetric to each other, where players choose between two actions C, D. If all players in each group choose the group action the individual payoffs are given by the the following matrix, where $\gamma, \lambda > 0$:

	C	D
C	$1\ ,\ 1$	$1-\gamma-\lambda\ ,\ 1+\gamma$
D	$1+\gamma\ ,\ 1-\gamma-\lambda$	$0\ ,\ 0$

Here γ is the transfer parameter, while λ measures inefficiency created by the transfer. If $1-\gamma-\lambda < 0$ the game is a prisoners dilemma (PD), while if $1-\gamma-\lambda > 0$ it is a game of chicken.

For this class of games we assume that the grid U starts at $\underline{u} < \min\{0, 1-\gamma-\lambda\}$. The other basic assumptions are as in Section 6.3. That is, we continue to assume that the grid has gaps of length $0 < d \leq \gamma, 1/2$ and that it does not contain the points $d/2, 1+d/2$. The leadership structure is unchanged (two group leaders and a common one, with $0 < v - v_3 < d$), and again U_0, U_1 denote the grid points closest to zero and 1.

In this class of games the probability of defection has an upper bound in any alliance constrained equilibrium. The following is proved in Appendix 2:

Theorem 8. *If ρ is an alliance constrained equilibrium of a transfer game then*

$$\rho(D,D) + \min\{\rho(C,D), \rho(D,C))\} \leq \frac{\gamma + d + v - v_3}{1+\gamma}$$

6.5. General Games with Group Leaders Who Can Talk

We now step back and ask: to which groups can leaders talk? The answer we propose is that leaders can approach groups to whom they can make credible offers - which are the groups who can punish them. Can a leader be punished by a group he does not care about? That is, if a leader's preference does not depend significantly on a given group's payoff, can he be punished by that group? If the answer is no, then group leaders, for example, can only talk to their own group (as we have assumed so far) and with this restricted message space the equilibrium payoff sets may be not be as large as one would hope. On the other hand it may also make sense to assume that group leaders can talk to other groups (see the discussion in Section 7.1). We now make this assumption, and show that it can make a sizable difference.

Assume there are two groups $k = \mathbf{1}, \mathbf{2}$ and just the two group leaders ℓ_1, ℓ_2 (ℓ_k with preferences of group k), but suppose that each of them can talk to both groups. As before, in case of a tie in utility bids to a group the own leader wins.

For the next result we assume that the grid interval length d is smaller than the least non-zero difference in payoffs to either group from any two action profiles. We let X_k be the action by group k that minmaxes group $-k$, where the maximization is subject to incentive compatibility. The following result holds:

Theorem 9. *Let $A = (A_1, A_2)$ be an efficient, incentive compatible profile where each group gets strictly more than the minmax payoff, and let U_1^A, U_2^A be the grid points closest to $u^1(A), u^2(A)$. Then profile A played for sure, with no leader punishment, is the outcome of the pure-strategy sequential equilibrium where the offers of leaders ℓ_1, ℓ_2 are respectively*

$$r^1 = (X_1, U_1^A - d), (A_2, U_2^A) \quad r^2 = (A_1, U_1^A), (X_2, U_2^A - d)$$

Proof. Note first that by definition an evaluator observing a leader's deviation believes the other leader is playing his equilibrium strategy. Consider leader 1's deviation possibilities. To get a higher payoff than $u^1(A)$, he has three options: keep control of group **2** alone, gain control of group **1** alone, or win control of both groups; and take the appropriate actions. To control group **2** alone he must bid (U_1, U_2) with $U_1 < U_1^A, U_2 \geq U_2^A$; and to increase his payoff he must prescribe to group **2** an action B_2 with $u^1(A_1, B_2) > u^1(A)$. But $U_1 < U_1^A$ implies evaluator 2 believes group **1** is following ℓ_2 hence playing A_1, and efficiency and the fine grid assumption imply $u^2(A_1, B_2) < u^2(A) - d \leq U_2^A - d/2 \leq U_2 - d/2$. This entails sure punishment by group 2.

What if $\ell_1$opted out of group **2** and took control of group **1** instead? His bid (U_1, U_2) should have $U_1 \geq U_1^A, U_2 < U_2^A$, plus some incentive compatible action prescription B_1 to group **1**. This implies group **2** playing the minmax X_2, evaluator 1 knowing it, and $u^1(B_1, X_2) < u^1(A) - d \leq U_1 - d/2$, that is no utility gain and sure punishment by group 1.

The only possibility left is to win both groups with a bid (U_1, U_2) where $U_1 \geq U_1^A, U_2 \geq U_2^A$, and have them play a profile B with $u^1(B) > u^1(A)$. But this implies that evaluator 2 observes offer (B_1, U_1), hence believes group **1** is following ℓ_1 and so is certain of the play of profile B. But by efficiency and the fine grid assumption $u^2(B) < u^2(A) - d \leq U_2 - d/2$, hence sure punishment of ℓ_1 by group **2**. □

To see what the result says in a familiar example reconsider the conformist's prisoner's dilemma of Example 6, with payoff matrix

	C	D
C	$1, 1$	$-\gamma, 1+\gamma$
D	$1+\gamma, -\gamma$	$0, 0$

In this case the theorem says that the efficient cooperative outcome (C, C) can be obtained for sure and without punishments in equilibrium with the following pair of offers, where U_0 and U_1 are the grid points closest to 0 and 1:

$$r^1 = (D, U_0), (C, U_1), \quad r^2 = (C, U_1), (D, U_0)$$

To see what is going on consider leader 1. In this game there is no point in inducing your group to cooperate unless the other is cooperating too; what the (C, U_1) part of offer

r^1 does is to "take the bull by the horns", convincing them to cooperate by asking them to do so under his leadership, effectively tying his own fate to theirs. The (D, U_0) part is to ensure that the other leader cannot achieve an even higher payoff for his own group without getting punished. In the end you "lead" your group to a high payoff by convincing the other group to follow your lead.

Still in the context of the prisoners dilemma notice that if the group leaders can only make offers to their own group the efficient outcome cannot be obtained despite the presence of a common leader. It is interesting that the common leader himself may stand to gain if the group leaders could talk to the other group.

On a different vein, one may observe that the "folk theorem" shown above might work with several groups where each group leader can talk to their own and adjacent group, but not if there are too many groups of the same size. Consider for simplicity the PD case. The generalization of the equilibrium found above has the group leader for group k talking to and winning group $k+1$, and the K-th group leader winning group 1. The potentially profitable deviation in this setting would involve a leader asking both his own group and the group he was meant to lead in equilibrium to defect while all the other groups cooperate. Then both groups take the offer, and make all the other groups pay. Whether this deviation works depends on how profitable it is for two groups to take advantage of everyone else. With three groups it is not so easy for two groups to take advantage of the third (two out of three groups deviating makes for a low payoff for all, so cooperation is an equilibrium), but with many groups a coalition of two groups is small and can reasonably get a high payoff by deviating while all the others are cooperating. We capture this observation formally in what follows.

Suppose there are K conformist groups that are engaged in the following *conformist class* game. For notational ease let group 1 be labeled as both 1 and $K+1$. Each group can choose either C or D. Suppose a fraction, ϕ of the groups choose D while the others choose C. Then each D group gets $\alpha(\phi)+\gamma(\phi)$ while each C group gets $\alpha(\phi)-\frac{\phi}{1-\phi}\gamma(\phi)-\lambda(\phi)$. We assume that α, γ and λ, respectively the per capita surplus from cooperation, transfer made to the defectors and the inefficiency caused by transfers, are all continuous functions of ϕ. In addition α and γ are non-increasing in ϕ and λ is non-decreasing and strictly positive. We further assume,

$$\alpha(0)+\gamma(0)>\alpha(0)>\alpha(1)+\gamma(1).$$

So for small enough K, $\phi=2/K$ is large enough to make $\alpha(\phi)+\gamma(\phi)<\alpha(0)$ but for K large enough $\phi=2/K$ will be so small that $\alpha(\phi)+\gamma(\phi)>\alpha(0)$, which makes cooperation not viable. Each group k has a group leader $\ell_k$who can make offers to groups k and $k+1$. Let $U_1(\phi)$ be the grid point closest to $\alpha(\phi)$ and d be smaller than the least non-zero difference in payoff to any group from any two action profiles.

Theorem 10. *There exists $\bar{K} > 0$ such that leader ℓ_k offering $(D, U_1 - d)$ to group k and (C, U_1) to group $k+1$ is a pure strategy sequential equilibrium if and only if $K \leq \bar{K}$. In this equilibrium all groups cooperate with certainty.*

Proof. Leader ℓ_k receives $\alpha(0)$ from the stated strategy profile. A profitable deviation must induce a higher payoff for Group k. There are only two possibilities. First, ℓ_k gains control over k by offering (D, U_1) while continuing to lead $k+1$ with the original offer. This, however, would result in sure punishment by Group $k+1$ since they would receive a payoff of $\alpha(1/K) - \frac{1}{K-1}\gamma(1/K) - \lambda(1/K)$ from such a deviation. Second, ℓ_k could offer (D, U_1) to k and also ensure that $k+1$ plays D by either offering them (D, U_1) or opting out by making a lower utility bid. Either way the deviation is unprofitable if and only if

$$\alpha(0) \geq \alpha(2/K) + \gamma(2/K).$$

Given our assumptions about the functions $\alpha$and γ there must exist some $\bar{K}$ such that the inequality above holds if and only if $K \leq \bar{K}$. □

7. Robustness

Besides their ubiquitous presence, the other relevant fact about leaders is that new ones may emerge. From an equilibrium perspective the question is then what happens if more leaders are added in a given situation. Which equilibria are robust to addition of new leaders? And, is competition among leaders good or bad for efficiency?

We first consider robustness in terms of the notion of *alliance proofness*. For a correlated strategy ρ we define the *utility set* $U^{\mathcal{K}}(\rho)$ for an alliance $\mathcal{K} \subseteq \{1, \ldots K\}$ as the set of utility vectors for members of the alliance corresponding to action profiles a that have positive probability in ρ. Let $\bar{v} = \max\{v^{\ell k} \mid \ell = 1, \ldots L, k \in \mathcal{K}^{\ell}\}$ be the highest leaders valence. We say that ρ is *ϵ-strongly blocked by* $\mathcal{K}$ if there exists a in-group $\overline{v}$- (pure) Nash equilibrium $a^{\mathcal{K}}$ that strictly Pareto dominates every point in $U^{\mathcal{K}}(\rho)$ by at least ϵ (for members of $\mathcal{K}$). We say that ρ is *ϵ-weakly alliance proof* if it is not ϵ-strongly blocked by any $\mathcal{K}^{\ell}$.

Theorem 11. *Every alliance constrained equilibrium is $\overline{v}$- weakly alliance proof.*

Proof. Suppose ρ is $\overline{v}$-strongly blocked for some $\mathcal{K}^{\ell}$. Let $a^{\mathcal{K}}$ be the blocking offer. Then ℓ can offer his alliance $a^{\mathcal{K}}$ giving them a utility vector $u^{\mathcal{K}}$. If d is the length of the grid interval above $u^{\mathcal{K}}$ he may bid up to $u^{\mathcal{K}} + d/2$ without being punished. On the other hand, regardless of which offer they accept, by assumption members of $\mathcal{K}$ get a utility $\tilde{u}^{\mathcal{K}}$ so a utility offer of at most $\tilde{u}^{\mathcal{K}} + d/2$. The attractiveness of this offer is at most $\tilde{u}^{\mathcal{K}} + d/2 + \overline{v}$. By assumption $\tilde{u}^{\mathcal{K}} + \overline{v} + d/2 < u^{\mathcal{K}} + d/2$ and the alternative offer by ℓ must be accepted. This offer will not be punished since it is $\overline{v}$-incentive compatible by assumption, and it makes the leader strictly better off. □

We now turn to robustness to the addition of leaders.

Definition 6. An alliance constrained equilibrium is *robust* if it remains an alliance constrained equilibrium when we add a leader identical to an existing leader but with a smaller valence than any existing leader.

We consider the robustness of equilibria encountered so far. No general result emerges. There are some good equilibria which are not robust and nasty ones which on the contrary seem impossible to disrupt. A comforting result is that the strongly symmetric equilibrium of Theorem 6 is robust. Recall that the game was the PD game with two group leaders talking to their own group and a common leader talking to both groups. If it is possible to sustain the equilibrium when entry takes place then it is certainly possible to do so when the evaluator of the new leader punishes the corresponding leader with probability 1 whenever he can do so, so we may restrict attention to this case.

Proposition 1. *The equilibrium of Theorem 6 in the conformist prisoners dilemma is robust.*

The proof of this is at the end of Web Appendix 2. Another robust equilibrium is not as nice. Consider the game-of-chicken transfer game, that is the class of Section 6.4 with $1 - \gamma - \lambda > 0$. The two classical asymmetric equilibria C, D and D, C are still equilibria with group leaders talking to own groups, the leader of the favored group offering D, U_1 (with U_1 still being the grid point near 1 that is near $1 + \gamma$).

Proposition 2. *Asymmetric equilibria in the game-of-chicken case are robust to addition of any type of new leader.*

Indeed, the leader of the group favored by the asymmetric equilibrium is winning own group with the high bid near $1 + \gamma$, so no type of leader can win this group and prescribe cooperation without getting punished for sure.

Other equilibria are not robust. Consider first the three-player Example *1.*

Proposition 3. *Suppose that the leader of the group is able to make offers to both groups. Equilibrium in Example* 1 *is then not robust to addition of a second group leader.*

Here the original leader will go for the $(5, 5, 5)$ payoff rather than the shadow mixed equilibrium. Now add a second leader for the group, able to make offers to the other group or not: the second leader goes for $(6, 6)$ for the group against the $(5, 5)$ equilibrium, so that equilibrium is not robust.

New group leaders seem disruptive, in particular they disrupt the efficient equilibria of Section 6.5. Let us focus on transfer games for simplicity, and maintain the assumption that $v - v_{new} < d < \gamma$ where v_{new} is the new leader's valence. To what kind of new leader addition is the C, C cooperative equilibrium robust? The answer is straightforward:

Proposition 4. *The cooperative equilibrium is robust to addition of a leader talking to group k if and only if he prefers that group to play C when the other group is playing C.*

The assertion is almost self evident: if he preferred group k playing D when $-k$ is playing C he could offer $(D, U_1 + d)$ to group k and opt out of $-k$. By so doing he would win k and disrupt the cooperative equilibrium. In other words we can add a new leader talking to k if he does not put too much weight to k's payoff. The cooperative equilibrium is robust to addition of common leaders but not to arrival of new group leaders. Emergence of new group leaders is definitely detrimental in this setting.

The result on the lack of robustness of the efficient equilibrium when more group leaders are added is luckily not the end of the story. Indeed, suppose there is a "pre-game" where groups can choose what sort of leadership structure to have. Efficiency requires that the group should choose to have a single leader who is susceptible to punishment by the other group - but has exactly the group preferences. Since the number of possible leaders is observed before bids are made, there is no disadvantage in having one leader. Basically by agreeing to have just one leader you give your leader the valuable possibility of making commitments without being undermined; by choosing someone who can be punished by the other group you give him the possibility of talking to the other group. Since he has your preferences you can trust him to do the best thing for your group. In other words groups may have the right incentives to develop efficient institutions.

An instance of efficient institutions in this respect is surprising enough - for the above story suggests that dynasties may be efficient in sustaining cooperation. Indeed a setting where a leader can be punished by opponent group - hence can credibly make offers to them - is the case of marrying off your ward to a competing family or kingdom. Consider then the efficient equilibrium implemented by marrying off your ward to the the neighboring country. The fact that new group leaders can disrupt this equilibrium corresponds to the rationale that the inter kingdom marriage strategy works as long as each kingdom is identified with a particular dynasty. A more democratic structure could not sustain such an equilibrium for more generations.

7.1. Which Groups Can Leaders Approach?

We initially introduced leaders as a concrete means of achieving group coordination. Our subsequent analysis makes it evident, however, that leadership structure plays a key role in the formation of coalitions and significantly affects efficiency and stability of equilibrium. So the key component of leadership structure in our model, namely the issue of which groups the leaders can make credible offers to, is basic. Since the incentive to be credible is given by punishment, we have to ask: by whom can leaders be effectively punished? Can a leader be punished by a group he does not care about? The answer clearly depends on how leaders are motivated. We briefly comment on this here.

The point is that leaders may have motivations different from that of the group they come from. In the exogenous group model this just means that they pick the incentive compatible outcome according to their own preference rather than the group preference. In the endogenous group model motivations play a much deeper role. The fact that leaders may be motivated by power or the desire to rule has important consequences. A leader who wants to rule the world can be punished by many groups - having his leadership repudiated by any group is costly to such a leader. A leader who wants to rule the world is not so likely to care only about parochial group interests. These people can be effectively punished by any group because what they care about first and foremost is to be leaders. Then all that is necessary for credibility in a group is that they know that even if you do not share their preferences you are hurt if they punish you because that hurts your leadership status.

If we were to think of this accountability as arising from some more fundamental preference, then perhaps we can ask if leaders who do not care about office (getting punished) ever become candidate leaders. In our model we explicitly rule these people out. But it may so happen that this caring about getting punished is simply a selection that happens in equilibrium. Self selection of leaders is a fundamental problem. One often thinks that those who do not want to be elected would be the best for office, but what we find is that "narcissism" helps after all since it makes the leader sensitive to punishment - which is key to credibility.

8. Conclusion

Our results cover two different areas. We study exogenously specified homogeneous collusive groups and argue strongly that the "right" notion of equilibrium is that of collusion constrained equilibrium by giving a number of different interpretations of that equilibrium. We then move on to endogenous and heterogeneous coalitions by adopting the approach of *ex post* evaluation which in the exogenous case gives rise to collusion constrained equilibrium. This leads to the notion of alliance constrained equilibrium. Our results here are more tentative and less complete but we think they represent a useful start.

In the exogenous homogeneous case we start from the observation that although Nash equilibrium does not account for collusion among subsets of players, when some subsets of players can be identified as potentially collusive groups, as is the case for example with political, ethnic or religious groups, collusion may influence group behavior. On collusion we adopt the obvious assumption that a group will collude on the within-group equilibrium which gives group members the highest utility when several equilibria exist. We find that this seemingly innocuous assumption disrupts existence of equilibrium in simple games. We show that the existence problem is due to a discontinuity of the equilibrium set, and propose a form of smoothing that overcomes the existence problem and results in a reasonable equilibrium concept which builds on the presumption that a group cannot be assumed to be able to play a particular within-group equilibrium with certainty when at that equilibrium the incentive constraints are satisfied with equality. This "tremble" implies that the group may put positive probability on actions which give group members lower utility but are strictly incentive compatible. We call the equilibrium "collusion constrained" because accounting for the possibility of collusive behavior on the part of some subsets of players constrains viable action profiles, with the consequence that in general collusion constrained equilibria payoffs are lower than straight Nash. We believe that the examples presented in the first half of the paper make a compelling case for collusion constrained equilibrium as the right starting point for analyzing exogenous groups (including dynamic models where people flow between exogenous groups based on economic incentives as in the Acemoglu (2001) farm lobby model), which in some sense is the case that Mancur Olson had in mind and is of key importance in a lot of existing political economy.

Focusing on group common actions we have then explored the role of group leaders as effective coordination devices, and have found that accountable leaders recommending actions would actually play recommendations constituting the collusion constrained equilibria found earlier.

In the second part of the paper we broaden leaders' role and consider more generally leaders competing for groups - in turn playing games between groups. Our model of leadership is coarse - it does not spell out the flow in information between leader and group as for example in Bolton, Brunnermeier and Veldkamp (2013) - but we think it does capture in a stark way the game that leaders play. Each candidate leader approaches different groups and each group is approached by different leaders. Leaders recommend actions and declare utility the group will achieve if they follow, and groups choose the leader who bids the highest utility but can punish him if the bid is unrealistically high. This simple structure has proven to yield a rich setting to study equilibria in games between groups.

In a series of examples, we show that equilibrium sets depend on the leadership structure, in particular on which groups the various leaders can approach and on the rules governing entry of new leaders - that is ultimately on institutions. For example, an efficient equilibrium we have studied is robust to addition of leaders caring for all groups but not to entry of additional leaders with preferences identical to those of a particular group. Institutions influence the ease of collusion within certain combination of groups and members and the ease with which new leaders emerge, and hence may have a significant impact on efficiency.

In summary we have studied situations where individual and group preferences are both at work, and each with a non-negligible weight bears upon the final outcome of a strategic game. The relative weight of the two forces is taken as given, for the balance between individual incentive constraints and group collusion ultimately depends on the exogenously given leaders' valences. Higher valences leave leaders' hands more free hence leave more room for group preferences, and vice versa. We have found that even for given valences the equilibria we have studied give interesting insights into group behavior and the impact of leaders.

References

Acemoglu, Daron (2001): "Inefficient Redistribution," *American Political Science Association*: 649-661.

Acemoglu, Daron and James A. Robinson (2000): "Why Did the West Extend the Franchise? Democracy, Inequality, and Growth in Historical Perspective, *Quarterly Journal of Economics* 115: 1167-1199.

Arrow, K. J. and G. Debreu (1954): "Existence of an equilibrium for a competitive economy," *Econometrica.*

Bernheim, B.Douglas, Bezalel Peleg, Michael D Whinston (1987): "Coalition-Proof Nash Equilibria I. Concepts," *Journal of Economic Theory* 42: 1-12.

Bolton, Patrick, Markus Brunnermeier and Laura Veldkamp (2013): "Leadership, Coordination and Corporate Culture", *Review of Economic Studies* 80: 512-537

Caillaud, Bernard and Philippe Jéhiel (1998): "Collusion in Auctions with Externalities," *RAND Journal of Economics* 29: 680-702.

Celik, G. (2009): "Mechanism design with collusive supervision," *Journal of Economic Theory.*

Che, Yeon-Koo and Jinwoo Kim (2006): "Robustly Collusion-Proof Implementation," *Econometrica* 74: 1063–1107.

Che, Yeon-Koo and Jinwoo Kim (2009): "Optimal Collusion-Proof Auctions," *Journal of Economic Theory.*

Cole, H. and E. C. Prescott (1997): "Valuation Equilibrium with Clubs," *Journal of Economic Theory* 74: 19-39.

Dewan, Torun and David Myatt (2008): "The Qualities of Leadership: Direction, Communication, and Obfuscation", *American Political Science Review* 102: 351-368

Dincecco, Mark, Giovanni Federico and Andrea Vindigni (2011): "Warfare, Taxation, and Political Change: Evidence from the Italian Risorgimento", *The Journal of Economic History* 71: 887-914

Ellickson, B., B. Grodal, S. Scotchmer and W. R. Zame (2001): "Clubs and the Market: Large Finite Economies," *Journal of Economic Theory* 101: 40-77.

Fudenberg, Drew and Jean Tirole (1991): Perfect Bayesian equilibrium and sequential equilibrium, *Journal of Economic Theory* 53: 236-260

Harsanyi, J. C. (1973): "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed-strategy Equilibrium Points," *International Journal of Game Theory* 2: 1-23.

Hermalin, Benjamin E. (1998): "Toward an Economic Theory of Leadership: Leading by Example", *American Economic Review* 88: 1188-1206

Hoffman, Philip T. and Jean-Laurent Rosenthal (2000): "Divided We Fall: The Political Economy of Warfare and Taxation", Mimeo, California Institute of Technology.

T. Hu, J. Kennan and N. Wallace (2009): "Coalition-Proof Trade and the Friedman Rule in the Lagos-Wright Model," *Journal of Political Economy.*

Laffont, Jean-Jacques and David Martimort (1997): "Collusion Under Asymmetric Information," *Econometrica* 65: 875-911.

McAfee, R. Preston and John McMillan (1992): "Bidding Rings," *American Economic Review* 82: 579-599.

McKelvey, R. D. and T. R.Palfrey (1995): "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior* 10: 6-38.

Olson, Mancur (1965): *The Logic of Collective Action: Public Goods and the Theory of Groups.*

Ray, Debraj and Rajiv Vohra (1999): "A Theory of Endogenous Coalition Structures," *Games and Economic Behavior* 26: 286-336.

Richerson, Peter J. and Robert Boyd (2001) "Commitment to Groups: A Tribal Instincts Hypothesis" in *Evolution and the Capacity for Commitment*, ed. Randolph M. Nesse, Russell Sage Foundations.

Rogers D. S., O. Deshpande and M. W. Feldman (2011): "The Spread of Inequality," PLoS ONE 6(9): e24683. doi:10.1371/journal.pone.0024683

Simon, L.K. and W.R. Zame (1990): "Discontinuous Games and Endogenous Sharing Rules," *Econometrica* 58: 861-872.

Tirole, J. (1986): "Hierarchies and Bureaucracies: on the Role of Collusion in Organizations" *Journal of Law, Economics and Organization* 2: 181-214.

Tirole, J. (1992): "Collusion and the Theory of Organizations," in J. J. Laffont, ed., *Advances in Economic Theory: Sixth World Congress*, Vol. II, Cambridge University Press, Cambridge.

Appendix 1: Correlation and Symmetry

We have so far supposed that the groups are homogeneous and that they choose only symmetric mixed strategies. We now wish to relax both of those assumptions. We first continue to assume that the group is homogeneous but allow a broader set of strategies. Then we show how the resulting model can be extended to heterogeneous groups in a way that is consistent with the homogeneous group model.

We have assumed that the strategies available to group k are a finite subset A^{kR} of symmetric mixed strategies, while the deviations available to individual members are the pure strategies A^k or the special strategy a_0^k meaning play the group mixed strategy a^k. Notice, however, that the assumption of symmetric mixed strategies is limiting. For example, if a group of two members is playing a hunter-gatherer game in which members choose between hunter and gatherer, and get 0 for agreeing, and the hunter gets 2 and the gatherer gets 1 if they specialize, the unique symmetric mixed equilibrium gives an expected utility to each member of 2/3 while a public randomization over the two asymmetric pure Nash equilibria gives an expected utility to each member of 3/2. In the game of chicken, for another example, there is a correlated equilibrium that gives both players more than any public randomization over Nash equilibria. It seems plausible that groups would choose to use correlating devices to achieve these superior results. This leads us to extend the model to include correlated strategies by each group.

In Section 3 we took the space of deviations to be $A^k \cup \{a_0^k\}$. By redefining A^{kR} and the space of deviations we can extend the model to incorporate correlated strategies in a straightforward way. First we take A^{kR} to be an arbitrary finite subset of symmetric *correlated* strategies for the group: that is, a probability distribution over profiles of individual actions. Then we define the space of deviations D^k to be maps $d^i : A^k \to A^k$ from pure actions to pure actions with the interpretation that $d^i(a^k)$ is the action chosen by member i when he is told to play a^k. Here the identity map plays exactly the role that a_0^k played in the original model. With this change all the existing results and definitions remain unchanged.

Extending the model to correlated strategies also enables us to incorporate asymmetries in a straightforward way. First, take A^{kR} to be an arbitrary finite subset of the correlated equilibria - not necessarily symmetric. We assume utility has the form $u^i(d^i, a^{k(i)}, a^{-k(i)})$ where $d^i \in D^{k(i)}$ and $a^{k(i)} \in A^{k(i)R}$, $a^{-k(i)} \in A^{-k(i)R}$ are no longer required to be symmetric, and individuals may no longer be homogeneous. The group is now assumed to have an exogenously specified objective of weighted sum of individual utility: $U^k(a^k, a^{-k}) = \sum_{i|k(i)=k} \omega_i u^i(a_0^k, a^k, a^{-k})$, and if we wish we may index the valences $v^i > 0$ by individual rather than by group. From a mathematical point of view, the only change needed to the existing model is that in the leadership version the evaluator must choose a vector of deviations $a^i|_{k(i)=k}$ and should equally weight the utility of each member of the group[15], while the leader should be punished if the evaluator chooses any deviation other than a_0^i on behalf of any group member. We refer to this notion as *asymmetric collusion constrained equilibrium.*

Given the asymmetric model, suppose the game is in fact symmetric - we would like to know that the new notion of equilibrium is consistent with the old notion. Suppose that

[15] Any strictly positive vector of weights is fine: we specify equal weights for definiteness. The point is that for the evaluator the optimal choice of each d is independent of the other choices.

the weights $\omega_i = 1$ and that the valences $v^i = v^{k(i)}$. Suppose also that for every correlated strategy $a^k \in A^{kR}$ the set A^{kR} also includes the uniform public randomization over all correlated strategies which permute the identities of the group members in a^k. In this case we say that A^{kR} *contains a symmetric model.* Then we can show that the new notion of asymmetric collusion constrained equilibrium is consistent with the old notion of symmetric collusion constrained equilibrium in the following sense:

Theorem 12. *Suppose that A^{kR}contains a symmetric model. Then there exists an asymmetric collusion constrained equilibrium $\tilde{\alpha}$ that is symmetric and is a collusion constrained equilibrium with respect to the subset of A^{kR} that is symmetric. Conversely if $\tilde{\alpha}$ is a collusion constrained equilibrium with respect to the subsets of A^{kR} that are symmetric then it is an asymmetric collusion constrained equilibrium.*

Proof. To show asymmetry implies symmetry, we construct the symmetric equilibrium from an arbitrary asymmetric equilibrium. Given a collusion constrained (or leadership) equilibrium - not necessarily symmetric - for each positive probability realization of the group public randomization device (or equivalently recommendation of the leader) we may replace the recommended profile a^k with the uniform public randomization over all permutations of the names of the group members, $\tilde{a}^k$. By assumption no other group cares about this, and since the incentive constraints are violated by no more than v^k at a^k for any group member $k(i) = k$ the same remains true for $\tilde{a}^k$. Moreover, $U^k(\tilde{a}^k, \alpha^{-k}) = U^k(a^k, \alpha^{-k})$ since each permutation of group member utilities yields exactly the same value. Hence $\tilde{a}^k$ is also an asymmetric collusion constrained equilibrium. Moreover, if $\tilde{a}^k$ gave less utility than some symmetric $\hat{a}^k$ that violates the incentive constraints by strictly less than v^k then so would a^k. Hence it is a symmetric collusion constrained equilibrium.

Now suppose that $\tilde{\alpha}$ is a collusion constrained equilibrium with respect to the subsets of A^{kR} that are symmetric and let $\tilde{a}^k$ be a positive probability realization of the group public randomization device. We have to show that there is no $\hat{a}^k \in A^{kR}$ that violates the incentive constraints by strictly less than v^k and has $U^k(\hat{a}^k, \tilde{\alpha}^{-k}) > U^k(\tilde{a}^k, \tilde{\alpha}^{-k})$. Suppose instead that there is such an $\hat{a}^k \in A^{kR}$. Consider the uniform randomization over permutations of group members of $\hat{a}^k$ and denote it by a^k. Then this also violates the incentive constraints by strictly less than v^k and has $U^k(a^k, \tilde{\alpha}^{-k}) = U^k(\hat{a}^k, \tilde{\alpha}^{-k}) > U^k(\tilde{a}^k, \tilde{\alpha}^{-k})$. But by construction a^k is symmetric and this then contradicts the fact that $\tilde{a}^k$ had positive probability in equilibrium. □

Appendix 2: Defection Probability Bound in Transfer Games

Theorem (Theorem 8 in text). *If ρ is an alliance constrained equilibrium of a transfer game then*

$$\rho(D, D) + \min\{\rho(C, D), \rho(D, C))\} \leq \frac{\gamma + d + v - v_3}{1 + \gamma}$$

We use the following

Lemma 2. *If an offer of u^{jk} is accepted from leader j by group k and $a^{jk} = C$ then $u^{jk} \leq U_1$.*

Proof. Since group k by accepting the offer gets at most 1 if a leader bids more than U_1 he is punished for certain and this is impossible in equilibrium. □

Proof of Theorem 8. Let U^{kk} be the highest positive probability bid of group leader k when offering D or $\underline{u}$ if he always offers C. If $U^{kk}+v<U_1+v_3$ for both group leaders then the common leader can offer $C, U_1, C, U_1$and guarantee that both groups cooperate, so avoid punishment and get a utility of 1. This is a strict improvement for the common leader unless the equilibrium is always cooperate, which is impossible.[16] Hence we conclude that for one group leader k we must have $U^{kk} \geq U_1-(v-v_3) \equiv \overline{U}_1$. Note that $U^{kk} \geq U_1-(v-v_3)$ and $(v-v_3)<d$ imply $U^{kk} \geq U_1$. Note, moreover, that group leader k does not have a bid $u>U^{kk}$ (by definition paired with C) accepted with positive probability, since from $v-v_3<d$ we get $U^{kk}>U_1-d$ hence $u>U^{kk} \geq U_1$ which implies by Lemma 2 that he cannot recommend C and have it accepted.

Define the unconditional equilibrium defection rates $\Upsilon_1=\rho(D, D)+\rho(D, C)$ and $\Upsilon_2=\rho(D, D)+\rho(C, D)$ and the corresponding defection rates $\Upsilon_k(\mathscr{E})$ conditional on event $\mathscr{E}$. If the common leader bids $u^{3k} \geq U^{kk}+v-v_3$ and there is a chance it is accepted then $a^{3k}=D$ by the lemma since $u^{3k} \geq U^{kk}+v-v_3>U_1$. Then $(1+\gamma)\left(1-\Upsilon_{-k}(u^{3k})\right)+d/2 \geq \overline{U}_1$ otherwise he is punished for sure. Since any such u^{3k} bid by the common leader in fact wins for certain we see that

$$\Upsilon_{-k}(u^{3k}) \leq 1-\frac{\overline{U}_1-d/2}{1+\gamma}.$$

Averaging over all the $u^{3k} \geq U^{kk}+v-v_3$ we see that this implies

$$\Upsilon_{-k}(u^{3k}>U^{kk}+v-v_3) \leq 1-\frac{\overline{U}_1-d/2}{1+\gamma}.$$

Moreover, the most obtainable for group leader k by bidding D, U^{kk} is no greater than $(1+\gamma)(1-\Upsilon_{-k}(u^{3k} \leq U^{kk}))$, or

$$(1+\gamma)(1-\Upsilon_{-k}(u^{3k} \leq U^{kk}+v-v_3)) \geq U^{kk}-d/2 \geq U_1-d/2$$

$$\Upsilon_{-k}(u^{3k} \leq U^{kk}+v-v_3)) \leq 1-\frac{U_1-d/2}{1+\gamma}.$$

Since the unconditional probability is an average of the two conditional probabilities and $\overline{U}_1<U_1$, we conclude that

$$\Upsilon_{-k} \leq 1-\frac{\overline{U}_1-d/2}{1+\gamma}.$$

[16] If there was such an equilibrium then it would necessarily be the case that group $-k$ cooperates regardless of the bid of group k: in this case group leader k can offer $D, 1+\gamma$, avoid punishment and be strictly better off.

Hence

$$\rho(D,D)+\min\{\rho(C,D),\rho(D,C))\}\leq 1-\frac{U_1+v_3-v-d/2}{1+\gamma}$$
$$=\frac{\gamma-v_3+v+1-U_1+d/2}{1+\gamma}\leq\frac{\gamma-v_3+v+d}{1+\gamma}$$

as indicated. □

Web Appendix 1: Example (Suburban Nightmare)

The role of this example is to highlight two issues: how collusion constrained equilibrium selects between different equilibria and how the inability to commit not to collude leads to inferior outcomes. In this example the notion of collusion constrained equilibrium captures how "keeping up with the Jones's" type preferences may be a problem not at the individual level but at the level of the family.

Three families live in a community. Each family i consists of a couple. Each member of each couple decides whether to work towards a shabby, casual or overdone lifestyle. The family members choose their actions simultaneously. The payoffs for a given family are as follows

	O	C	S
O	-2,-2	-2,2	-2,1
C	2,-2	2,2	-2,4
S	1,-2	4,-2	1,1

The payoffs above capture the fact that each family member prefers a casual lifestyle to a shabby one, while hating an overdone lifestyle. The hatred (-2) for working towards an overdone lifestyle is unaffected by what one's partner is up to. However, working towards a casual lifestyle is made very difficult if one's partner shirks and works for shabby instead. This is captured by the payoff $(4,-2)$ where the row player gets the benefit of a comfortable lifestyle while maintaining shabby habits. The hardworking partner in the family is of course not amused with the -2.

The lifestyle choice of each member of the family contributes to a lifestyle image for the family as a whole. The image of Overdone requires both family members to choose overdone. The image of casual requires at least one family member to choose casual. All other within-family choices lead to an image of shabby. Each family member gets additional utility or disutility depending on where they stand on their community lifestyle ranking. For some absurd reason Overdone is considered better than Casual which is better than Shabby.

$$O \succ_{RANK} C \succ_{RANK} S$$

Being ranked first *alone* brings an additional utility of 5. Being last *alone* adds -4. Not being first but not being alone at that rank brings -2. All other ranks result in no additional utility or disutility for a family. For example, if all families are at the same rank then there is no additional utility or disutility. So the payoff to a particular family member is simply the one derived from the matrix above determined by choices made in the family plus the payoff received from the ranking of the family, determined by choices made across families.

Each family behaves collusively. The unique pure strategy collusion constrained equilibrium of this game is every member of each family choosing *overdone.*

Treating families as individuals yields four predictions, (O, O, O) with payoff $(-2, -2, -2)$ and the 3 different permutations of (C, C, O) with payoff (relevant permutation of) $(0, 0, 3)$.

Instead, treating a family simply as two individuals without the ability to collude, brings us back to the standard notion of Nash equilibrium. No strategy profile that results in a family image profile of (C, C, O) can be a Nash equilibrium. To see this observe that one of the members of family 1 must be choosing C. Then the other member is strictly better off choosing S. Doesn't change the image (and therefore rank) but brings a higher personal payoff of 4. Now if a member of family 1 is choosing $S$and the family has an image of C then the other family member must be choosing C and getting a personal payoff of -2 plus the common family payoff of -2 (not first but not alone either). So a total of -4. Deviating to S affects the family image and its rank. It makes family 1 come last with a common payoff of -4. But at least now the family member who deviated gets an additional personal payoff of 1. So in net, -3.

But families playing (C, C, C) with individual payoffs higher than in the collusion constrained equilibrium can arise from a Nash equilibrium. In each family one member chooses C while the other chooses S. The payoffs are $((4, -2), (4, -2), (4, -2))$ The family member who gets the sucker's payoff in each family does not prefer to deviate to S because that would reduce the rank of their family to *last alone.* This would bring a net payoff to this member of $-4 + 1 = -3$.

Web Appendix 2: Analysis of the PD Example

We analyze strongly symmetric equilibrium in the additively separable prisoner's dilemma game

	C	D
C	$1 , 1$	$-\gamma , 1+\gamma$
D	$1+\gamma , -\gamma$	$0 , 0$

with individual preferences for conformity.

Lemma 3. *(i) Provided expected utility is not exactly at the midpoint of a grid interval, predicted utility is the unique grid point in the length-d neighborhood of expected utility. (ii) In any equilibrium, if a leader bid exceeds expected utility by more than $d/2$ and is accepted the leader is punished for sure.*

Proof. (i) Recall that if expected utility is in the lower [resp. upper] half of a grid interval then predicted utility is the lower [resp. upper] bound of the interval. Now let $\mathbb{E}u^k$ be expected utility and $u \in U$ be the unique grid point $\mathbb{E}u^k - d/2 < u < \mathbb{E}u^k + d/2$; then either $u < \mathbb{E}u^k < u + d/2$ or $u > \mathbb{E}u^k > u - d/2$, and by what just recalled in both cases predicted utility is u. (ii) Suppose $u_1, u_2, u_3 \in U$ are consecutive grid points, that expected utility $\mathbb{E}u^k \in [u_1, u_2]$ and bid is $u > \mathbb{E}u^k + d/2$. If $\mathbb{E}u^k < u_1 + d/2$ then predicted utility $u^j = u_1$ while $u \geq u_2$. If $\mathbb{E}u^k \geq u_1 + d/2$ then $u^j \leq u_2$ while $u \geq u_3$. So in both cases $u > u^j$ and the leader is punished. □

We also recall that if a group accepts a bid of u and expected utility is exactly $u - d/2$ the evaluator can punish the leader with any probability for he is indifferent between choosing u and $u - d$ as predicted utility. Recall that $u^{\ell k}$ is the utility offered to group k by leader ℓ.

Definition. *[Definition 5 in the text]* A *strongly symmetric equilibrium* is one in which both group leaders play the same strategy, offer only D and bid greater than or equal to U_0 and in which the offers made by the common leader are of the form C, u, C, u.

For the remainder of the Appendix equilibrium refers to strongly symmetric equilibrium.

Lemma 4. *In an equilibrium the common leader's accepted offers involve bids of no more than $1 + d/2$, that is they never exceed U_1.*

Proof. If the common leader has an accepted offer of C, u^{31}, C, u^{32} each group gets at most 1 so if the common leader bids more than $1 + d/2$ to either group he is punished for certain (Lemma 3(ii)). □

Recall that $R(u)$ is the probability with which a group leader bids less than or equal to u and $Q(u)$ is the probability with which the common leader bids less than or equal to u. We now prove Theorem 6 in the text - it follows from Propositions 5 and 6 below.

Lemma 5. *If group leaders never recommend C and do not bid $u_k < U_0$ then if the common leader has a profitable deviation to D, D or D, C he also has one of the form C, C.*

Proof. Deviating to D, u^{31}, D, u^{32} results in both groups playing D for sure, so is exactly the same as deviating to $C, \underline{u}, C, \underline{u}$. If with positive probability the common leader has a bid D, u^{31}, C, u^{32} accepted by group 2 then the actual utility received by group 2 when that bid is accepted must be non-positive so to avoid certain punishment, the group leader must be bidding $u^{32} \leq d/2$ which by the generic assumption means $u^{32} \leq U_0$. Since the group leaders are not bidding less than U_0 we conclude in fact that the bid D, u^{31}, C, u^{32} is rejected by group 2, so that it results in both group playing D for sure. Hence deviating to D, u^{31}, C, u^{32} is exactly the same as deviating to $C, \underline{u}, C, \underline{u}$, which also results in both groups playing D for sure. □

Lemma 6. *A strongly symmetric equilibrium exists.*

Proof. If we restrict the strategy space of the group leaders to D and the common leader to C, u, C, u then by Nash (1951) we know that an equilibrium exists in which the two group leaders play the same mixed strategy over their bids. This is due to the game being symmetric with respect to the two group leaders.

Now we show that in the restricted game there is a symmetric equilibrium in which $R(U_0 - d) = 0$, that is, group leaders do not bid below U_0. By the generic assumption $U_0 < d/2$ so the group leader cannot be punished for such an offer, because by recommending D his group cannot get less than 0. This implies that the group leader switching all bids $u^{kk} < d/2$ to U_0 weakly dominates the original plan since for the group leader D strictly dominates C. Moreover, if there is positive probability of an offer by the common leader with with $u \leq d/2$ the group leader does strictly better by switching, so in equilibrium this is not the case. In other words all offers by the common leader with $u \leq d/2$ are rejected by both group leaders with probability 1. Hence a group leader shifting all bids $u_k < d/2$ to U_0 does not change the play path, nor does it matter to the other group leader. We just need to check that we have not introduced an incentive for the common leader to underbid the group leader: however a bid by the group leader with $u \leq d/2$ will not be punished, and an underbid against the new strategy loses for sure, so is the same as bidding $\underline{u}$ against the old strategy and that was not an improvement for the common leader.

Next we show that the group leaders have no incentive to offer C. Consider the bid C, u^{kk}, and observe that if D, u^{kk} is bid instead it wins exactly when C, u^{kk} would have. Hence if C, u^{kk} has positive probability of winning the leader does strictly better by bidding D, u^{kk}.

Finally we show that the common leader has no incentive to deviate to the strategies we have excluded. We showed in Lemma 5 that we need only consider deviations by the common leader of the form C, C. So we need only show that if there is a profitable deviation of the form u^{31}, u^{32} then there is one of the form u, u. Assume without loss of generality that $u^{31} < u^{32}$. Expected utility of group k if it accepts this bid is $\overline{u}_k = R(u^{3,-k} - d) - \gamma(1 - R(u^{3,-k} - d)) = (1 + \gamma)R(u^{3,-k} - d) - \gamma$. Since the offer is off the equilibrium path, we may assume indifferent evaluators punish with probability 1, so it must be that $u^{3k} < (1 + \gamma)R(u^{3,-k} - d) - \gamma + d/2$ so that common leader is not punished. His utility is then $R(u^{31} - d)R(u^{32} - d) + R(u^{31} - d)(1 - R(u^{32} - d))/2 + (1 - R(u^{31} - d))R(u^{32} - d)/2 = (R(u^{31} - d) + R(u^{32} - d))/2$.

We also have $u^{32} < (1+\gamma)R(u^{31} - d) - \gamma + d/2 \leq (1+\gamma)R(u^{32} - d) - \gamma + d/2$. Hence the offer u^{32}, u^{32} is also unpunished and gets utility $R(u^{32} - d) \geq (R(u^{31} - d) + R(u_2^{32} - d))/2$, so if the deviation u^{31}, u^{32} is profitable, so is u^{32}, u^{32} . □

Now we analyze strongly symmetric equilibrium.

Notation. We now let $\overline{u}_k$ be the highest bid that leader k plays with positive probability. Also, set $q(u) = Q(u) - Q(u - d)$ as the probability that the common leader bids exactly u and for $Q(u) > 0$ define $H(u) \equiv \sum_{u_3 \leq u} q(u_3)P(u_3 - d)/Q(u)$, which is the probability that a group cooperates conditional on the common leader bidding less than or equal to u.

Lemma 7. *The equilibrium expected utility of a group when its leader bids U_1 is $(1+\gamma)H(U_1)$.*

Proof. Since bidding U_1 beats the common leader for sure, the group will defect for sure, so will receive utility equal to the probability that the other group leader is outbid by the common leader $H(U_1)$ times the payoff from playing D against C which is $1+\gamma$. □

Lemma 8. *There can be no equilibrium bid that wins with probability 1 and is punished with probability 0; moreover, in equilibrium all leaders have a positive probability of winning.*

Proof. If the group leaders never win the bidding then they get 1. In this case if a group leader were to bid U_1 he would win for sure, get $1+\gamma$ for sure, and since $1+\gamma > 1$ not get punished. This contradicts the fact that the group leaders only get 1 in equilibrium.

If the common leader could make a bid that wins with probability 1 and does not get punished he would receive 1 for certain, so his equilibrium utility would have to be 1. This implies that the group leaders never win the bidding, which we just showed is impossible.

Suppose that the common leader never wins the bidding. Then the actual utility received by all the players is 0, so that no leader bids more than $d/2$ with positive probability. But then the common leader can bid $U_1 + d$ and win with probability 1 without being punished, which we just showed is impossible.

Finally, if a group leader could make a bid that wins with probability 1 and does not get punished he would get $(1+\gamma)H(U_1)$. On the other hand, any bid $u < U_1$ would receive strictly less utility. Hence the common leader never wins the bidding, which we just showed is impossible. □

Lemma 9. *In equilibrium $\overline{u}_3 = \overline{u}_1 = U_1$, the unique grid point $1 - d/2 < U_1 < 1 + d/2$.*

Proof. $\overline{u}_3 \leq U_1$ is Lemma 4. If $\overline{u}_1 < 1 - d/2$ then the common leader can bid $\overline{u}_1 + d$, win for sure and not be punished so $\bar{u}_1 \geq U_1$, hence $\bar{u}_1 \geq \bar{u}_3$. Neither can $\overline{u}_1 > \overline{u}_3$ be true. Since $\bar{u}_1$ wins with probability 1 we know by Lemma 8 that it must be punished with positive probability. Since this probability cannot be one either it requires the evaluator to be indifferent about punishment. So the expected payoff from following the action must equal the bid minus $d/2$. By bidding $\bar{u}_3$ instead the group leader continues to win with probability 1, generates the same outcome but avoids punishment. Finally if $\overline{u}_3 > \overline{u}_1$ then it must be that $1 + d/2 = \overline{u}_3$ and $1 - d/2 = \overline{u}_1$ which is ruled out by the generic assumption on the grid.. □

Proposition 5. The equilibrium probability a group cooperates, that is $H(U_1)$, is given by

$$\frac{1-d}{1+\gamma} \leq H(U_1) = \frac{U_1 - d/2}{1+\gamma} \leq \frac{1}{1+\gamma}$$

Proof. When a group leader bids $U_1 = \overline{u}_3$ he wins for sure and group utility is $(1+\gamma)H(U_1)$ (Lemma 7). Since he wins for sure he must be punished with positive probability, and since he plays this bid with positive probability we must have $(1+\gamma)H(U_1) = U_1 - d/2$ so that the evaluator is indifferent to punishing him. Inequalities follow from $1-d/2 < U_1 < 1+d/2$. □

Lemma 10. *If u has positive probability of acceptance for the common leader in equilibrium then*

$$R(u-d) \geq \frac{u+\gamma-d/2}{1+\gamma}.$$

and if it has positive probability of acceptance for the group leaders in equilibrium then

$$H(u) \geq \frac{u-d/2}{1+\gamma}.$$

Proof. These are the conditions that the leaders not be punished with probability 1. The first follows from the fact that if common leader bid is accepted the group expected utility is $R(u-d)-\gamma(1-R(u-d)) = (1+\gamma)R(u-d)-\gamma$ (and Lemma 3). The second is analogous. □

Definition 7. $u \in U$ is a *positive point* for a leader if in equilibrium the leader plays it with positive probability and is punished with positive probability.

Lemma 11. *If u is a positive point for the common leader then the group leaders play $u-d$ with positive probability; if u is a positive point for the group leaders then the common leader plays u with positive probability. At a positive point for the common leader*

$$R(u-d) = \frac{u+\gamma-d/2}{1+\gamma}.$$

At a positive point for a group leader

$$H(u) = \frac{u-d/2}{1+\gamma}.$$

The point $U_1 = \overline{u}_1 = \overline{u}_3$ is a positive point for both types of leaders.

Proof. The first part just says that a leader should not be able to lower his bid, leave chance of winning unchanged and reduce probability of being punished. Equalities follow from the fact that if u is a positive point then expected utility must be exactly $u-d/2$. The point $\overline{u}_1 = \overline{u}_3$ is played by both types of leaders with positive probability by definition, as it is the largest such point. If the group leader plays $\overline{u}_1$.he wins with probability 1 hence by Lemma 8 he must be punished with positive probability. If $\bar{u}_3$ were unpunished when accepted then the common leader should play it with probability 1 since the common leader does strictly better by playing $\overline{u}_3$ then any other bid; this cannot happen in equilibrium. □

Lemma 12. *The equilibrium probability of common leader bids satisfies*

$$Q(u-d) = \frac{R(u-d)-H(u)}{R(u-d)-H(u-d)}Q(u).$$

If u is a positive point for both leaders then and $u-d$ is a positive point for the group leaders then

$$Q(u-d) = \frac{\gamma}{\gamma+d}Q(u).$$

Proof. We have $H(u)Q(u) = \sum_{u_3 \le u} q(u_3)R(u_3 - d)$, so

$$\begin{aligned} H(u)Q(u) - H(u-d)Q(u-d) &= \sum_{u_3 \le u} q(u_3)R(u_3 - d) - \sum_{u_3 \le u-d} q(u_3)R(u_3 - d) \\ &= q(u)R(u-d) = (Q(u) - Q(u-d))R(u-d) \end{aligned}$$

At positive points we use the values of R and H given above. $\square$

Proposition 6. *There is a unique strongly symmetric equilibrium. In this equilibrium no leader bids below U_0. For $U_0 \le u < U_1$*

$$R(u) = \frac{u + \gamma + d/2}{1 + \gamma}$$

with $R(U_1) = 1$ and for $U_0 < u \le U_1$

$$Q(u) = (\gamma/(\gamma + d))^{(U_1 - u)/d}$$

with

$$Q(U_0) = (\gamma/(\gamma + d))^{(U_1 - U_0 - d)/d} \left(\frac{\gamma}{U_0 + \gamma + d/2} \right).$$

Neither leader is punished for bidding U_0 and both leaders are punished for each higher bid.

Proof. Recall that u not being a positive point means that u is played with probability zero or is played with positive probability and punished with probability zero if accepted by at least one group. For each type of leader let $\hat{u}_\ell$ be the largest point below $\overline{u}_\ell = U_1$ that is not a positive point for that leader, and let $\hat{u} = \max\{\hat{u}_1, \hat{u}_3\}$. Define

$$\hat{\ell} = \begin{cases} 1 & \text{if } \hat{u}_1 \ge \hat{u}_3 \\ 3 & \text{if } \hat{u}_1 < \hat{u}_3 \end{cases}$$

First $\hat{\ell}$ plays $\hat{u}$ with strictly positive probability and is not punished for doing so; neither player has a positive point at or below $\hat{u}$. This follows from Lemma 11: in case $\hat{\ell} = 1$, since $\hat{u} + d$ is a positive point for the common leader then the group leaders play $\hat{u}$ with positive probability; in case $\hat{\ell} = 3$, since $\hat{u}$ is a positive point for the group leaders then the common leader plays $\hat{u}$ with positive probability. It follows directly that $\hat{u}$ is not punished for $\hat{\ell}$. So $\hat{\ell}$ cannot have a positive point below $\hat{u}$: it would be strictly better to switch to $\hat{u}$.

Suppose $\hat{\ell} = 1$. We first show that $\hat{u}$ must also be accepted with positive probability. Suppose not: then the common leader does not bid anything below his positive point $\hat{u} + d$. By the definition of H we get,

$$H(\hat{u} + d) = R(\hat{u})$$

Now since $\hat{u} + d$ is a positive point for both the common leader and the group leader it must be that,

$$\frac{\hat{u} + d - d/2}{1 + \gamma} = H(\hat{u} + d) = R(\hat{u}) = \frac{\hat{u} + d + \gamma - d/2}{1 + \gamma}$$

a contradiction.

So if $\hat{\ell} = 1$ then any positive probability offer by 3 at or below $\hat{u}$ must lose with probability 1: if not it beats some positive probability offer of the group leader, who should switch to $\hat{u}$ not be punished and strictly increase the probability of winning; this also implies that 3 does not have a positive point at or below $\hat{u}$: if so 3 could lose equally well by bidding less and avoid punishment.

As a result $H(\hat{u}) = 0$, which along with Lemma 10 gives

$$0 = H(\hat{u}) \geq \frac{\hat{u} - d/2}{1+\gamma}$$

This implies that $\hat{u} \leq d/2$, which because the equilibrium is simple implies $\hat{u} = U_0$. So if $\hat{\ell} = 1$ it must be that $\hat{u} = U_0$.

Next, suppose $\hat{\ell} = 3$. Firstly it must be that the bid $\hat{u}$ wins with positive probability for the common leader. Since otherwise $R(\hat{u} - d) = 0$, implying that $H(\hat{u}) = 0$. But since $\hat{u}$ is a positive point for the group leaders we get

$$H(\hat{u}) = \frac{\hat{u} - d/2}{1+\gamma}$$

a contradiction.

Then any positive probability offer by 1 strictly below $\hat{u}$ must lose with probability 1: if not it beats some positive probability offer of the common leader, who should switch to $\hat{u}$ not be punished and strictly increase the probability of winning; this also implies that 1 does not have a positive point strictly below $\hat{u}$ since again, it would be better to lose by bidding less and avoiding punishment.

From this construction since the group leaders lose with probability 1 bidding strictly below $\hat{u}$, let $\tilde{u}$ be the highest group leader bid with positive probability strictly below $\hat{u}$. By definition there are no positive probability offers by the group leader above $\tilde{u}$ and at or below $\hat{u} - d$ so $R(\tilde{u}) = R(\hat{u} - d)$. Moreover, for the same reason for $\tilde{u} < u_3 \leq \hat{u}$ we have $R(u_3 - d) = R(\hat{u} - d)$. For $u_3 \leq \tilde{u}$ we must have $q(u_3) = 0$ since otherwise the group leaders would not lose with probability 1. By definition $H(\hat{u})Q(\hat{u}) \equiv \sum_{u_3 \leq \hat{u}} q(u_3)R(u_3 - d) = \sum_{\tilde{u} < u_3 \leq \hat{u}} q(u_3)R(u_3 - d) = \sum_{\tilde{u} < u_3 \leq \hat{u}} q(u_3)R(\hat{u} - d) = Q(\hat{u})R(\tilde{u})$ or $H(\hat{u}) = R(\tilde{u}) = R(\hat{u} - d)$. Since $\hat{u}$ is a positive point for the group leader, we know from Lemma 11 that

$$R(\hat{u} - d) = H(\hat{u}) = \frac{\hat{u} - d/2}{1+\gamma}.$$

Since $\hat{u}$ has positive probability of acceptance for the common leader by Lemma 10

$$R(\hat{u} - d) \geq \frac{\hat{u} + \gamma - d/2}{1+\gamma}.$$

Hence we have the inequality

$$\frac{\hat{u} - d/2}{1+\gamma} = R(\hat{u} - d) \geq \frac{\hat{u} + \gamma - d/2}{1+\gamma}$$

which is impossible. Hence we conclude that there is no positive probability bid by group leaders strictly below $\hat{u}$. But this implies that 3 loses with probability 1 by bidding $\hat{u}$, contradicting our earlier finding that if $\hat{l} = 3$ then the bid $\hat{u}$ wins with positive probability for the common leader. So it cannot be that $\hat{\ell} = 3$.

At the bottom we have

$$H(U_0 + d) = \frac{U_0 + d/2}{1+\gamma} = \sum_{u_3 \leq U_0+d} q(u_3)R(u_3 - d)/Q(U_0 + d) = q(U_0 + d)R(U_0)/Q(U_0 + d)$$

along with

$$R(U_0) = \frac{U_0 + \gamma + d/2}{1+\gamma}$$

and from Lemma 12

$$Q(U_0 + d) = (\gamma/(\gamma + d))^{(U_1 - U_0 - d)/d}.$$

Then

$$\frac{U_0 + d/2}{1+\gamma}(\gamma/(\gamma + d))^{(U_1 - U_0 - d)/d} = q(U_0 + d)\frac{U_0 + \gamma + d/2}{1+\gamma}$$

from which

$$q(U_0 + d) = \frac{U_0 + d/2}{U_0 + \gamma + d/2}(\gamma/(\gamma + d))^{(U_1 - U_0 - d)/d}$$

Since the equilibrium is simple, nobody bids below U_0 so

$$Q(U_0) = q(U_0) = Q(U_0 + d) - q(U_0 - d) = (\gamma/(\gamma + d))^{(U_1 - U_0 - d)/d}\left(1 - \frac{U_0 + d/2}{U_0 + \gamma + d/2}\right)$$

Since the common leader must be indifferent across the bids she mixes over, her expected payoff from any of those bids must be 0, since that is the payoff she gets from making the lowest bid, namely U_0. Making a bid greater than U_1 would result in punishment with certainty, a worse outcome. Bidding less than U_0 is not a profitable deviation either since it loses with certainty.

For the group leader, expected payoff from bidding U_0 is

$$q(U_0)(0) + \sum_{U_0 < u \leq U_1} q(u)[R(u - d)(1) + (1 - R(u - d))(-\gamma)]$$

A lower bid (that must therefore always lose) would instead give a payoff of

$$q(U_0)(-\gamma) + \sum_{U_0 < u \leq U_1} q(u)[R(u - d)(1) + (1 - R(u - d))(-\gamma)]$$

It therefore does not profit the group leader to make a bid lower than U_0. Finally the group leader wouldn't bid greater than U_1 since it would result in certain punishment. □

For the continuum limit we have

Theorem. *[Theorem 7 in the text] The limit of the unique strongly symmetric equilibrium as $d \to 0$ is given by*

$$R(u) = \frac{u+\gamma}{1+\gamma}$$

with $R(1) = 1$ and for $0 \le u \le 1$

$$Q(u) = e^{(u-1)/\gamma}.$$

Proof. Analyzing $(\gamma/(\gamma+d))^{(U_1-u)/d}$ by taking logs we have $(U_1-u)[\log\gamma - \log(\gamma+d))]/d \to (u-U_1)/\gamma$. □

We lastly prove robustness of this equilibrium:

Proposition. *[Proposition 1 in text] The equilibrium of Theorem 6 in the conformist prisoners dilemma is robust.*

Proof. Suppose a new leader makes a bid of $V_1 \ge V_2$ to group 1 and group 2 respectively with recommendation a_1, a_2, where he is assumed to have a smaller valence than either the group or common leader. For existing leaders, we let u be the common leader bid and u_1, u_2 the group leaders bids.

Notice that if a leader who can talk to just one group has a winning bid, so does a leader who can talk to both groups, since he can always intentionally lose the bidding in either group. The following is the table of possible outcomes where $*$ means that the new offer does not win in either group.

	$u_1 < V_1, u_2 < V_2$	$u_1 < V_1, u_2 \ge V_2$	$u_1 \ge V_1, u_2 < V_2$	$u_1 \ge V_1, u_2 \ge V_2$
$u < V_2$	a_1, a_2	a_1, D	D, a_2	$*$
$V_2 \le u < V_1$	a_1, C	$a_1, C-D$	D, C	$*$
$u \ge V_1$	$*$	$*$	$*$	$*$

If group **2** accepts the offer, the only two possibilities are a_1, a_2 and D, a_2 so that if D is proposed to group **1** then conditional on **2** accepting the offer, group **2** either gets 0 or $-\gamma$, and in either case honesty - in the sense of avoiding punishment for sure - compels the new leader to lose the bid. Hence the new leader in this case is bidding only to group **1** and is bidding D. Conditional on the bid winning in group **1** the fact that he beat the group **1** leader contains no information about group **2** play, so he faces exactly the same distribution for group **2** as if the group **1** leader bid $V_1 - d$, which is to say the actual utility received by group **1** is $V_1 - d/2$ so it seems the new leader can't bid D to group **1** and win the bidding with positive probability. Hence the only possible bids (that won't lose with probability 1 or be punished with probability 1) by the new leader are C, D and C, C.

If $V_1 = V_2$ then by the same reasoning the new leader can't bid D to group **2**. So the only possibilities are: $a_1 = a_2 = C$; $V_1 > V_2 > U_0$,$a_1 = C$,$a_2 = D$; $V_2 = U_0$

Suppose **2** accepts the bid, then conditional utility of the group is $R(V_1-d)u(C,a_2) + [1 - R(V_1-d)]u(D,a_2)$ that is $R(V_1-d)[u(C,a_2) - u(D,a_2)] + u(D,a_2)$. If $a_2 = C$ then $R(V_1-d)[u(C,C)-u(D,C)] + u(D,C) = R(V_1-d)[1+\gamma] - \gamma = \frac{V_1+\gamma-d/2}{1+\gamma}[1+\gamma] - \gamma = V_1 - d/2$ which is indifferent so the evaluator can punish with probability 1, so this case is impossible. If on the other hand $a_2 = D$ then $R(V_1-d)[u(C,D)-u(D,D)] + u(D,D) = (1+\gamma)\frac{V_1+\gamma-d/2}{1+\gamma} = V_1 - d/2 + \gamma > V_2 - d/2 + \gamma$.

So the remaining cases are: $V_1 > V_2$,$a_1 = C$,$a_2 = D$; and $V_2 = U_0$. Observe that if the leader cannot make a credible bid to group **1** when $V_2 = U_0$ then he cannot be better off by also telling group **2** to defect and winning some of the time. So we are left to consider the case $V_2 = U_0$. Since D is dominant for group **1** if there is a credible bid for group **1**, it must be when the leader tells group **1** to play D. So: can he tell group **1** to play D and promise them some utility V_1 when he makes a bid only to that group? The negative answer to this question is contained in the first part of this proof. The conclusion is then that any bid that the new leader can make that is accepted with positive probability results in punishment with probability 1 (assigning the evaluator to punish with probability 1 when indifferent). So a new leader will not enter, the equilibrium is robust. □