

Collusion Constrained Equilibrium[☆]

Rohan Dutta¹, David K. Levine², Salvatore Modica³

Abstract

We study collusion within groups in non-cooperative games. The primitives are the preferences of the players, their assignment to non-overlapping groups and the goals of the groups. Our notion of collusion is that a group coordinates the play of its members among different incentive compatible plans to best achieve its goals. Unfortunately, equilibria that meet this requirement need not exist. We instead introduce the weaker notion of *collusion constrained equilibrium*. This allows groups to put positive probability on alternatives that are suboptimal for the group in certain razor's edge cases where the set of incentive compatible plans changes discontinuously. These collusion constrained equilibria exist and are a subset of the correlated equilibria of the underlying game. We examine four perturbations of the underlying game. In each case we show that equilibria in which groups choose the best alternative exist and that limits of these equilibria lead to collusion constrained equilibria. We also show that for a sufficiently broad class of perturbations every collusion constrained equilibrium arises as such a limit. We give an application to a voter participation game showing how collusion constraints may be socially costly.

Keywords: Collusion, Organization, Group

[☆]First Version: June 2, 2013. We are especially grateful to Marco Mariotti, Andrea Mattozzi, Kirill Pogorelskiy, Joel Sobel, Chris Tyson, seminar participants at Bocconi, Glasgow, Naples, Padua, QMUL, the Games meetings in Maastricht and a number of anonymous referees. We are grateful to the EIEF 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)

¹Department of Economics, McGill University

²Department of Economics, EUI and WUSTL

³Università di Palermo

1. Introduction

As the literature on collective action (for example Olson (1965)) has emphasized, groups often behave collusively while the preferences of individual group members limit the possible collusive arrangements that a group can enter into. Neither individual rationality - ignoring collusion - nor group rationality - ignoring individual incentives - provides a satisfactory theory of interaction between groups. We study what happens when collusive groups face internal incentive constraints. Our starting point is that of a standard finite simultaneous move non-cooperative game. We suppose that players are exogenously partitioned into groups and that these groups have well-defined objectives. Given the play of the other groups there may be several Nash equilibria within a particular group (within-group equilibria). We model collusion within that group by supposing that the group will agree to choose the within-group equilibrium that best satisfies its objectives.

The idea of choosing a best outcome for a group subject to incentive constraints has not received a great deal of theoretical attention but is important in applications. It has been used in the study of trading economics, for example, by Hu, Kennan and Wallace (2009). In industrial organization Fershtman and Judd (1986) study a duopoly where owners employ managers. Kopel and Loffler (2012) use a similar setting to explore asymmetries. Balasubramanian and Bhardwaj (2004) study a duopoly where manufacturing and marketing managers bargain with each other. In other settings the group could be a group of bidders in an auction as in McAfee and McMillan (1992) and Caillaud and Jéhiel (1998), or it might consist of a supervisor and agent in the Principal/Supervisor/Agent model of Tirole (1986).⁴ In political economy Levine and Modica (2016)'s model of peer pressure and its application to the role of political parties in elections by Levine and Mattozzi (2016) use the same notion of collusion. In mechanism design a related idea is that within a mechanism a particular group must not wish to recontract in an incentive compatible way. A theoretical study along these lines is Myerson (1982).⁵

The key problem that we address is that *strict collusion constrained equilibria* in which groups simultaneously try to satisfy their goals subject to incentive constraints do not

⁴See also the more general literature on hierarchical models discussed in Tirole (1992) or Celik (2009). For other types of mechanisms see Laffont and Martimort (1997) and Martimort and Moreira (2010). Most of these papers study a single collusive group. One exception is Che and Kim (2009) who allow multiple groups they refer to as cartels. In the theory of clubs, such as Cole and Prescott (1997) and Ellickson et al (2001), implicitly collusion takes place within (many) clubs - but the clubs interact in a market rather than a game environment.

⁵Myerson also observes that there is an existence problem and introduces the notion of quasi-equilibrium to which our collusion constrained equilibrium is closely connected. This link is explored in greater detail below. We should emphasize that while our notion of equilibrium and existence result are similar to Myerson's, unlike Myerson, our primary focus is on examining what is captured by the notion of equilibrium and consequently on whether it makes sense.

generally exist. That is, if groups take the actions of other groups as given and choose the best Nash equilibrium within the group no equilibrium may exist. Example 1 illustrates. For this reason applied theorists have generally either avoided imposing individual incentive constraints on group actions or else invented *ad hoc* solutions to the existence problem.⁶ We show that the existence problem is due to the discontinuity of the within-group equilibrium correspondence and show how it can be overcome by allowing, under certain razor's edge conditions, randomizations by groups between alternatives to which they are not indifferent. This leads to what we call *collusion constrained equilibrium*. This is a special type of correlated equilibrium of the underlying non-cooperative game.

Our key goal is to motivate our definition of collusion constrained equilibrium. We argue that it is useful because it correctly captures several different types of small influences that might not be convenient to model explicitly. Specifically we consider three perturbations of the underlying model. We first consider models in which there is slight randomness in group beliefs about the play of other groups. This provides a formal version of the informal arguments we use to motivate the definition. We then consider models in which groups may overcome incentive constraints at a substantial enforcement cost - that is, group members are allowed take suboptimal non-equilibrium actions but the group must bear enforcement costs to induce them to do so.⁷ For both of these perturbations strict collusion constrained equilibria exist⁸ - in particular randomization occurs only when there is indifference - and as the perturbation vanishes the equilibria of the perturbed games converge to collusion constrained equilibria of the underlying game. Finally, we explore the Nash program of motivating a cooperative concept as a limit of non-cooperative games. Specifically, we consider a model in which there is a non-cooperative meta-game played between “leaders” and “evaluators” of groups and in which leaders have a slight valence.⁹ If we call the leaders

⁶For example. in the collusion in auction literature as stated in Harrington (2008), it is assumed that non-colluding firms will act the same in an industry with a cartel as they would without a cartel.

⁷Collusion constrained equilibrium allows groups to put positive probability on alternatives that are suboptimal for the group in certain cases. The model of costly enforcement does not allow this, but instead allows individual group members to play suboptimal but nearly optimal strategies.

⁸There is a certain irony here: using enforcement to overcome incentive constraints is quite natural in a principal-agent setting. This result shows that even if enforcement is quite costly the existence problem noted by Myerson (1982) in the principal-agent setting goes away. On the other hand if enforcement is quite costly it is natural to work with the limiting case where enforcement is not possible and our results show that collusion constrained - or quasi - equilibrium correctly capture what happens in that case.

⁹A related class of models, for example Hermalin (1998), Dewan and Myatt (2008) and Bolton, Brunnermeier and Veldkamp (2013), examine leadership in which a group benefits from its members coordinating their actions in the presence of imperfect information about the environment. In this literature, however, there is no game between groups - the problem is how to exploit the information being acquired by leader and group members in the group interest. For example, Bolton, Brunnermeier and Veldkamp (2013) find that the leader should not put too much weight on the information coming from followers (what they call “resoluteness” of the leader).

“principals” this formulation is the closest to the models used in mechanism design. In the leader/evaluator game perfect Bayesian equilibria exist and as the valence approaches zero once again the equilibrium play path converges to a collusion constrained equilibria of the underlying game.

These upper hemicontinuity results with respect to the three perturbations show that the set of collusion constrained equilibria is “big enough” in the sense of containing the limits of equilibria of several interesting perturbed models. The second key question we address is whether the set of collusion constrained equilibria is “too big” in the sense that perhaps not all collusion constrained equilibria arise as such limits - indeed, we could capture all relevant limits trivially by defining everything to be an equilibrium. Could there be a stronger notion of equilibrium that still captures the relevant limits? For any particular perturbation the answer is yes: we show in a simple example that limits from perturbed games lead to strict refinements - that is, subsets - of collusion constrained equilibria - albeit different refinements depending on which perturbation we consider. Is it also the case that the set of collusion constrained equilibria is too big because some collusion constrained equilibria do not arise as any limit from interesting perturbed games? In our final theoretical result we show that this is not the case. We consider a combination of two perturbations: a belief and an enforcement cost perturbation, and to eliminate non-generic preferences allow also a perturbation to the group objective. Once again in these perturbed games strict collusion constrained equilibria exist and converge to collusion constrained equilibria of the underlying game. However, for this broader class of perturbations we have the converse as well: all collusion constrained equilibria of the underlying game arise as such limits. Hence our key conclusion: the set of collusion constrained equilibria is “exactly the right size,” being characterized as the set of limit points of strict collusion constrained equilibria for this broad yet relevant class of perturbations.

In our theory incentive constraints play a key role. In applied work the presence of incentive constraints within groups has often been ignored. For example political economists and economic historians often treat competing groups as single individuals: it is as if the group has an unaccountable leader who makes binding decisions for the group. In Acemoglu and Robinson (2000)’s theory of the extension of the franchise there are two groups, the elites and the masses, who act without incentive constraints. Similarly in the current literature on the role of taxation by the monarchy leading to more democratic institutions the game typically involves a monarch and a group (the elite).¹⁰ In our leader/evaluator perturbation we also assume that the group decision is made by a single leader, but we add to the game

¹⁰Hoffman 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).

evaluators who punish the leader for violating incentive constraints. We focus on strategic interaction between groups and a central element of our model is accountability, in that a leader whose recommendations are not endorsed by the group will be punished.

We should emphasize that there is an important territory between ignoring incentive constraints entirely and requiring as we do that they always be satisfied. An important example that we study explicitly is the possibility that incentive constraints can be overcome - for example through an enforcement mechanism - albeit at some cost. Here we can view “no incentive constraints” as “no cost of enforcement” on the one extreme and “incentive constraints must always be satisfied” as “very high cost of enforcement” on the other. One result that we establish is to give conditions on costly enforcement such that strict collusion constrained equilibria do exist. More broadly our contribution is oriented towards applications where incentive constraints cannot easily be overcome.

One branch of the game theory literature that is closely connected to the ideas we develop here is the literature that uses non-cooperative methods to analyze cooperative games. There, however, the emphasis has been on the endogenous formation of coalitions - generally in the absence of incentive constraints. The Ray and Vohra (1997) model of coalition formation contains in it a theory of how exogenously given groups play a game among themselves. The present paper may be viewed as a simple building block for their theory of coalition formation. With exogenous groups an equilibrium in their model requires groups to play strategy profiles that given the behavior of the other groups cannot be Pareto improved. Notice that there are no incentive constraints.¹¹ Our analysis is at a lower level: we look for a reasonable solution when a partition of players is exogenously given and groups interact with each other strategically, keeping within-group incentive compatibility constraints. There is also an extensive literature that describes the game by means of a characteristic function and involves proposals and bargaining. 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. However the details, goals, and analysis are quite different: while we analyze deviations by a fixed group for a given partition, Bernheim, Peleg and Whinston (1987) consider any coalitional deviation and analyze which deviations are credible.¹²

To make the theory more concrete we study an example based on the voter participation

¹¹Haeringer (2004) points out that the assumption of quasi concave utility is insufficient in guaranteeing equilibrium existence in the Ray and Vohra (1997) setup with exogenous groups, unless the groups can play within-group correlated strategies. The non-existence problem in our setup is of an entirely different nature and, in particular, is independent of whether or not groups can play correlated strategies.

¹²Because coalitions are not pre-defined the existence problem for coalition proof Nash equilibrium is more akin to the problem of an empty core than to the continuity problems discussed here.

model of Palfrey and Rosenthal (1985) and Levine and Mattozzi (2016). We consider two parties - one larger than the other - voting over a transfer payment and we depart slightly from the standard model by assuming that ties are costly. In this setting we find all the Nash equilibria, all the collusion constrained equilibria, and all the equilibria in which the groups have a costless enforcement technology. We study how the equilibria compare as the stakes are increased. The main findings for this game are the following. For small stakes nobody votes. For larger stakes in Nash equilibrium it is always possible for the small party to win. If the stakes are large enough in collusion constrained and costless enforcement equilibrium the large party preempts the small and wins the election. For intermediate stakes strict collusion constrained equilibria do not exist, but collusion constrained equilibria do. For most parameter configurations the collusion constrained equilibria are more favorable for the large party than Nash equilibrium, less favorable than costless enforcement equilibrium, and less efficient than either.

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 groups subject to incentive constraints. The basic idea we explore in this paper is that if, given the play of other groups, there is more than one within-group equilibrium, then a collusive 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) action profile that gives them the most utility. However, in the following game there is no equilibrium that satisfies this prescription.

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 both cooperate (play C)

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

If player 3 plays D the resulting payoffs are as follows, where notice that players 1 and 2 are then in a coordination game:

	<i>C</i>	<i>D</i>
<i>C</i>	10, 10, 0	0, 8, 5
<i>D</i>	8, 0, 5	2, 2, 5

Let α^i denote the probability with which player i plays C . We examine the set of within-group equilibria for players 1 and 2 given the strategy α^3 of player 3. The payoff matrix for those two players is

	<i>C</i>	<i>D</i>
<i>C</i>	$6 + 4(1 - \alpha^3), 6 + 4(1 - \alpha^3)$	0, 8
<i>D</i>	8, 0	2, 2

so that as α^3 starts at 1 the two players face a prisoner's dilemma game with a unique within-group Nash equilibrium at D, D , and as α^3 decreases the payoff to cooperation is increasing until at $\alpha^3 = 1/2$ the game becomes a coordination game and the set of within-group equilibria changes discontinuously with a second pure strategy within-group equilibrium at C, C ; for $\alpha^3 < 1/2$ there is an additional symmetric strictly mixed within-group equilibrium in which $\alpha^1 = \alpha^2 = 1/2(1 - \alpha^3)$.

How should the group of player 1 and player 2 collude given the play of player 3? Let us suppose that the group objective satisfies the Pareto criterion. If $\alpha^3 > 1/2$ they have no choice: there is only one within-group equilibrium at D, D . For $\alpha^3 \leq 1/2$ they each get $6 + 4(1 - \alpha^3)$ at the C, C within-group equilibrium, 2 at the D, D within-group equilibrium, and strictly less than $6 + 4(1 - \alpha^3)$ at the strictly mixed within-group 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 which is the best within-group 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 $\alpha^3 = 0$ which is not larger than $1/2$; when the group plays C, C player 3 should play C and so $\alpha^3 = 1$ which 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 within-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, for as α^3 is increased slightly above .5 the C, C within-group equilibrium abruptly vanishes. The key idea of this paper is that this discontinuity is a shortcoming of the model rather than an intrinsic feature of the underlying group behavior. To motivate 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: 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 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 within-group 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 in the first case agree that $\alpha^3 \leq 0.5$ with 90% probability and mistakenly agree that $\alpha^3 > 0.5$ with 10% probability 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 this limit: 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 within-group equilibrium.¹³ That is, when the incentive constraint holds exactly we do not assume that the group can choose their most preferred within-group equilibrium, but instead we assume that there is an endogenously determined probability that they will choose that within-group equilibrium. In this case optimality for player 3 requires her to be indifferent between C and D , so in the “collusion constrained” equilibrium we propose the group will mix 50-50 between C, C and D, D ; player 3 mixes 50-50 between C and D .

The import of collusion constraints can be seen by comparing what happens in this game without collusion. This game has three Nash equilibria: one at D, D, D , one in which

¹³This is similar to Simon and Zame (1990)'s endogenous choice of sharing rules.

3 plays D and 1 and 2 mix 50-50 between C and D , and a fully mixed one.¹⁴ In the first the group members each get 2 in the second 5 and in the third 6.25. By contrast in the unique collusion constrained equilibrium the group members each get 5. Moreover in the completely mixed Nash equilibrium player 3 gets 2.5 exactly as in the unique collusion constrained equilibrium. Why do not the group members get together and promise player 3 not to collude and instead coordinate on the completely mixed Nash equilibrium? They will be better off and player 3 is indifferent. The problem is that by saying that player 1 and 2 form a group we mean that they cannot credibly commit not to collude. If such an agreement was reached with player 3 as soon as the meeting was over players 1 and 2 would convene a second meeting among themselves and agree that rather than mixing they will play C, C . Anticipating this player 3 would never make the original agreement. It would be convenient for public policy if lobbying groups - such as bankers and farmers - could credibly commit not to collude among themselves. Unfortunately this is not the world we live in, hence the need to consider collusion constraints.

Remark. Discontinuity and non-existence is not an artifact of restricting attention to within-group Nash equilibrium. The same issue arises if we assume that players 1 and 2 can use correlated strategies. When the game is a prisoner's dilemma, that is, $\alpha^3 > 1/2$ then strict dominance implies that the unique within-group Nash equilibrium is also the unique within-group correlated equilibrium. When $\alpha^3 \leq 1/2$ the within-group correlated equilibrium set is indeed larger than the within-group Nash equilibrium set (containing at the very least the public randomizations over the within-group Nash equilibria), but these within-group correlated equilibria are all inferior for players 1 and 2 to C, C and 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.

The bad behavior of the best-equilibrium correspondence is related to some of the earliest work on competitive equilibrium. Arrow and Debreu (1954) showed that the best choice from a constraint set is well-behaved when the constraint set is lower hemicontinuous. If it is, then the maximum theorem can be applied to show that the argmax is a continuous correspondence.¹⁵ However, neither the Nash nor correlated equilibrium correspondence used as a constraint set is lower hemicontinuous, and - as we have seen - the best-equilibrium correspondence can then fail to be continuous.¹⁶

¹⁴The game is analyzed in Web Appendix 2.

¹⁵This is the approach used by Myerson (1982) to prove the existence of quasi-equilibrium.

¹⁶A specific scenario in which the discontinuity and the non-existence problem goes away requires all players belonging to the same group to have identical ordinal preferences over all outcomes.

3. Collusion Constrained Equilibrium

3.1. The Environment

We now introduce our formal model of collusive groups that pursue their own interest subject to within-group individual incentive constraints. The membership in these groups is exogenously given and the ability of a group to collude is independent of actions taken by players outside of the group. We emphasize that we use the word collusion in the limited meaning that the group can choose a within-group equilibrium to its liking. The goals of the group - like those of individuals - are exogenously specified: we do not consider the possibility of conflict within the group over goals.

Our basic setting is that of a standard normal form game. There are players $i = 1, 2, \dots, I$; player i chooses actions from a finite set $a^i \in A^i$ and receives utility $u^i(a^i, a^{-i})$. On top of this standard normal form game we have the structure of groups $k = 1, 2, \dots, K$. There is a fixed assignment of players to groups $i \mapsto k(i)$. Notice that each player is assigned to exactly one group and that the assignment is fixed and exogenous. We use $a^k \in A^k$ to denote (pure) profiles of actions within group k and a to denote the profile of actions over all players. Like individuals, groups have well-defined objectives given by a payoff function $v^k(a^k, a^{-k})$.¹⁷

We assume that groups can make plans independently from other groups. We take this to mean that each group k has an independent *group randomizing device* the realization of which is known to all group members but not to players who are not group members. One implication of this is that the play of group k appears from the perspective of other groups to be a correlated strategy - a probability distribution $\rho^k \in R^k$ over pure action profiles A^k . In addition to the group randomizing device the individual players in a group can randomize, so that by using the group randomizing device the group can randomly choose a profile of mixed strategies for group members. We let $\alpha^k \in \mathcal{A}^k$ represent such a profile, albeit we take $\mathcal{A}^k \subseteq R^k$ so that rather than regarding α^k as a profile of mixed strategies we choose to regard it as the generated distribution over pure strategy profiles A^k . Formally if α^i denote probability distributions over A^i then $\alpha^k[a^k] \equiv \prod_{k(i)=k} \alpha^i[a^i]$. If the group mixes over a subset $B^k \subseteq \mathcal{A}^k$ using the group randomizing device the result is in the convex hull of B^k which we write as $H(B^k)$.

¹⁷Notice that we are not restricting the group objective function. Depending on the application, some group objectives might be more natural than others. For example we might have $v^k(a^k, a^{-k}) = \sum_{i|k(i)=k} \beta^i u^i(a^k, a^{-k})$ for some positive utility weights $\beta^i > 0$. This implies on the one hand a preference for Pareto efficient plans, but also agreement on the welfare weights. In the special case of a group with just one individual such a group objective function might be especially compelling. On the other hand, considerations of fairness might cause a group of more than one individual to prefer a Pareto inferior plan.

Players choose deviations $d^i \in D^i = A^i \cup \{0\}$ where the deviation $d^i = 0$ means “mix according to the group plan.”¹⁸ Individual utility functions then give rise to a function

$$U^i(d^i, \alpha^k, a^{-k}) = \begin{cases} \sum_{a^k} u^i(a^i, a^{k-i}, a^{-k}) \alpha^k[a^k] & d^i = 0 \\ \sum_{a^k} u^i(d^i, a^{k-i}, a^{-k}) \alpha^k[a^k] & d^i \neq 0 \end{cases} .$$

It is convenient also to have a function that summarizes the degree of incentive incompatibility of a group plan. Noting that the randomizations of groups are independent of one another, for $\alpha^k \in \mathcal{A}^k, \rho^{-k} \in R^{-k}$ we define

$$G^k(\alpha^k, \rho^{-k}) = \max_{i|k(i)=k, d^i \in D^i} \sum_{a^{-k}} \left(U^i(d^i, \alpha^k, a^{-k}) - U^i(0, \alpha^k, a^{-k}) \right) \prod_{j \neq k} \rho^j[a^j] \geq 0$$

which represents the greatest expected gain to any member of group k from deviating from the plan α^k given the play of the other groups. The condition for group incentive compatibility is simply $G^k(\alpha^k, \rho^{-k}) = 0$.

The key properties of the model are embodied in $G^k(\alpha^k, \rho^{-k})$ and

$$v^k(\alpha^k, \rho^{-k}) = \sum_a v^k(a^k, a^{-k}) \alpha^k[a^k] \prod_{j \neq k} \rho^j[a^j]$$

Both functions are continuous in (α^k, ρ^{-k}) and it follows from the standard existence theorem for Nash equilibrium in finite games that for every ρ^{-k} there exists an α^k such that $G^k(\alpha^k, \rho^{-k}) = 0$. These properties together with \mathcal{A}^k being a closed subset of R^k are the only properties that are used in the remainder of the paper.¹⁹ For example, we might wish to allow groups to choose among within group correlated equilibria rather than within group Nash equilibria. In this case we can take \mathcal{A}^k to be all correlated strategies by group k .²⁰ Alternatively if we thought that homogeneous group might be restricted to anonymous play we can take \mathcal{A}^k to be the mixed strategy of a representative individual.

¹⁸Note that since we are dealing here with ordinary mixed strategies there is no need to consider deviations conditional on the outcome of the individual randomizing device. See also footnote 20

¹⁹It would not in general be appropriate to assume \mathcal{A}^k convex for the following reason. We want public randomizations over incentive compatible plays. But a distribution over profiles which is a correlated equilibrium (hence incentive compatible) with respect to some correlating device is not necessarily generated by public randomization over incentive compatible profiles. For example, a group which has no correlating devices available except public randomization cannot achieve the usual $(1/3, 1/3, 1/3)$ correlated equilibrium in the game of chicken without violating incentive compatibility, because that distribution is obtainable only through the public randomization that puts weight $1/3$ on the three pure strategy profiles - which are not all incentive compatible. However a convex \mathcal{A}^k containing the pure profiles would also contain $(1/3, 1/3, 1/3)$. We must thus dispense with a convexity assumption on \mathcal{A}^k to properly account for incentive compatibility within groups.

²⁰In this case we must broaden the definition of a deviation to be a contingent deviation: $d^i : A^i \rightarrow A^i$ reflecting a choice of how to play contingent on a particular recommendation.

We are now in position to give a comparison of our setup with that of Myerson (1982). Myerson adds to the model a finite set of types for each player. This in itself does not change anything: our actions can easily be the finite set of maps from types to individual decisions. However Myerson's types are reported to a group coordinator (the principal) who can then make recommendations to individual group members. We do not allow this - so that our model corresponds to Myerson's model where there is a single type of each group member. In this sense our model is a specialization of Myerson. However, because the principal can make private recommendations, in Myerson the space $\mathcal{A}^k = R^k$ the space of all correlated strategies for the group. As indicated our model is consistent with this possibility but it is not our base model and we do not require this. In the setting of groups rather than agents organized by principals it is not particularly natural to allow such a broad set of group strategies, nor is it particularly natural in the applied game settings described in the introduction. For example, if the groups are groups of voters it makes sense that they might coordinate their play by communicating a degree of enthusiasm for voting that day (randomizing as a group), but less sense that they would communicate their individual difficulty of voting that day to a central coordinator who would then make individualized voting recommendations.

The bottom line of this comparison is that our model in which $\mathcal{A}^k = R^k$ (and the appropriate set of deviations are used, see footnote 20) coincides with the Myerson model in which there is a single type. As we indicate below - in this case his notion of quasi-equilibrium is exactly our definition of collusion constrained equilibrium.

3.2. Equilibrium

We first give a formal definition of the notion of strict collusion constrained equilibrium. As we have already shown that these may not exist, we then go on to consider collusion constrained equilibrium.

Recall that $G^k(\alpha^k, \rho^{-k})$ measures the greatest gain in utility to any group member of deviating from the plan α^k . The greatest incentive compatible group utility is given by

$$V^k(\rho^{-k}) = \max_{\alpha^k \in \mathcal{A}^k \mid G^k(\alpha^k, \rho^{-k})=0} v^k(\alpha^k, \rho^{-k})$$

For the solutions to the maximization problem we say:

Definition 1. The *group best response set* $B^k(\rho^{-k})$ is the set of plans α^k satisfying $G^k(\alpha^k, \rho^{-k}) = 0$ and $v^k(\alpha^k, \rho^{-k}) = V^k(\rho^{-k})$.

Note that $B^k(\rho^{-k})$ is closed. We can then define

Definition 2. $\rho \in R$ is a *strict collusion constrained equilibrium* if $\rho^k \in H[B^k(\rho^{-k})]$ for all k .

As these may not exist we now give our definition of *collusion constrained equilibrium*. We adopt the motivation given in Myerson (1982) for his notion of quasi-equilibrium.²¹ Recall that in the proposed collusion constrained equilibrium of our example the 3rd player was randomizing 50-50 and that as a consequence it was a within-group equilibrium for the group to either both cooperate or both defect. However cooperation is not a safe option in the sense that a small perturbation in beliefs can cause it to fail to be incentive compatible. Hence group members might be concerned that after an agreement is reached some small change in beliefs would lead members to violate of the agreement. On the other hand defection is safe in the sense that if such an agreement is reached no small change in beliefs would lead any group member to wish to violate the agreement. Let us first define the ϵ -worst best utility for group k for beliefs near ρ^{-k} as

$$V_\epsilon^k(\rho^{-k}) = \inf_{|\sigma^{-k} - \rho^{-k}| < \epsilon} V^k(\sigma^{-k}).$$

Observe that this is non-increasing in ϵ , so we may take the limit and define $V_S^k(\rho^{-k}) = \lim_{\epsilon \rightarrow 0} V_\epsilon^k(\rho^{-k})$ as the *group safety utility*. Our basic premise is that there would be no reason for the group to choose a plan which gives less group utility than the group safety utility. For incentive compatible plans yielding higher utility we are agnostic - perhaps the group can reach agreement on such plans, perhaps not. This leads us to define:

Definition 3. The *shadow*²² *response set* $B_S^k(\rho^{-k})$ is the set of plans α^k that satisfy $G^k(\alpha^k, \rho^{-k}) = 0$ and $v^k(\alpha^k, \rho^{-k}) \geq V_S^k(\rho^{-k})$.

Like $B^k(\rho^{-k})$ we have $B_S^k(\rho^{-k})$ closed. Note that since $V_S^k(\rho^{-k}) \leq V^k(\rho^{-k})$ we have $B_S^k(\rho^{-k}) \supseteq B^k(\rho^{-k})$. We know from example 1 that $B^k(\rho^{-k})$ may fail to be upper hemicontinuous since a sequence of incentive compatible best plans may converge to a plan that is not best. We show in the Appendix that by contrast the correspondence $B_S^k(\rho^{-k})$ must be upper hemicontinuous. The key intuition is that the group safety level $V_S^k(\rho^{-k})$ can jump down but not up so that a sequence of safe plans converge to a safe plan.²³ Because $B_S^k(\rho^{-k})$ is upper hemicontinuous $B_S^k(\rho^{-k}) = B^k(\rho^{-k})$ implies that $B^k(\rho^{-k})$ is also upper hemicontinuous at ρ^{-k} and we say that ρ^{-k} is a *regular point* for group k . Otherwise we say that ρ^{-k} is a *critical point* for group k .

²¹As indicated above in the case where the two models coincide the definition of quasi-equilibrium and collusion constrained equilibrium coincide as well.

²²The set $B_S^k(\rho^{-k})$ is a kind of shadow of nearby best within-group equilibria.

²³Basically this solves the problem of discontinuity by forcing Reny (1999)'s better reply security condition. Bich and Laraki (2016) demonstration that Reny solutions are Nash is similar to the fact here that when the shadow best response set is the same as the group best response set then collusion constrained equilibria are strict collusion constrained equilibria.

Our premise is that the group will place weight only on incentive compatible plans which provide at least the group safety utility, that is, on $B_S^k(\rho^{-k})$ so we adopt the following definition.

Definition 4. $\rho \in R$ is a *collusion constrained equilibrium* if $\rho^k \in H[B_S^k(\rho^{-k})]$ for all k .

The key to collusion constrained equilibrium is that we allow plans in $B_S^k(\rho^{-k})$ not merely in $B^k(\rho^{-k})$. If in a collusion constrained equilibrium $\rho^k \notin H[B^k(\rho^{-k})]$ we say that group k engages in *shadow mixing*. This means that the group puts positive probability on within-group equilibria in $B_S^k(\rho^{-k}) \setminus B^k(\rho^{-k})$, that are not the best possible.

Our example above shows that shadow mixing may be necessary in equilibrium, as we spell out next.

Example. [*Example 1 revisited*] In the example we take $k(1) = k(2) = \mathbf{1}, k(3) = \mathbf{2}$. In this and all subsequent use of this example we take group utility to be defined by equal welfare weights on individual utility functions $v^1(a^1, a^2) = u^1(a^1, a^2) + u^2(a^1, a^2)$ and $v^2(a^1, a^2) = u^3(a^1, a^2)$.

To apply the definition of collusion constrained equilibrium we first compute for group $k = \mathbf{1}$ the best utility $V^1(\rho^2)$ where since there is one player ρ^2 may be identified with α^3 . For $\alpha^3 \leq 1/2$ we know that the best within-group equilibrium for group $k = \mathbf{1}$ is C, C with corresponding group utility $V^1(\rho^2) = 12 + 8(1 - \alpha^3)$, while for $\alpha^3 > 1/2$ the only within-group equilibrium is D, D with group utility $V^1(\rho^2) = 4$. For $\alpha^3 \neq 1/2$ we have $V_S^1(\rho^2) = V^1(\rho^2)$, and the shadow response and best response sets are the same: C, C for $\alpha^3 < 1/2$ and D, D for $\alpha^3 > 1/2$. At $\alpha^3 = 1/2$ the worst best utility for nearby beliefs are those for $\alpha^3 > 1/2$ giving a group utility of 4, whence the set of incentive compatible plans that give at least this utility are the within-group equilibria C, C and D, D , that is $B_S^1(\rho^2) = \{(C, C), (D, D)\}$. For the group $k = \mathbf{2}$ consisting solely of individual 3 the shadow response set is just the usual best response set.

Clearly there is no equilibrium with $\alpha^3 \neq 1/2$. On the other hand when $\alpha^3 = 1/2$ the group can shadow-mix 50-50 between C, C and D, D , leaving player 3 indifferent between C and D ; so this is a collusion constrained equilibrium. We conclude that there is a unique collusion constrained equilibria with ρ^1 a 50-50 mixture over $\{(C, C), (D, D)\}$ and ρ^2 a 50-50 mixture over $\{C, D\}$.

As the example shows collusion constrained equilibrium may require that the group sometimes agree to plans that are “unsafe.” Whether this makes sense is not clear - one of our main tasks in the remainder of the paper will be to establish whether it does indeed make sense.

It should be apparent that collusion constrained equilibria use as correlating devices only the private randomization device available to each player and the group randomiza-

tion device. We refer to correlated equilibria of the underlying game that use only these randomizing devices as *group correlated equilibria*.²⁴ Formally, let $\mathcal{B}^k(\rho^{-k})$ be the set of plans α^k satisfying $G^k(\alpha^k, \rho^{-k}) = 0$. then a group correlated equilibrium is a $\rho \in R$ such that $\rho^k \in H[\mathcal{B}^k(\rho^{-k})]$ for all k .

Theorem 1. *Collusion constrained equilibria exist and are a subset of the group correlated equilibria of the underlying game.*

The theorem is proved in the Appendix.²⁵ It makes clear the sense in which collusion constrained equilibria are constrained: there are many group correlated equilibria, but the ones that are interesting from the point of view of collusion are those in which groups are constrained to play in their shadow response sets.

4. Three Model Perturbations

We now study how collusion constrained equilibrium arises as a limit of equilibria in perturbed models. The key point is that equilibria in the perturbed models will be strict: groups make best choices and there is no shadow mixing. There is no issue of the group sometimes sacrificing utility for safety and sometimes not. Nor is there an issue of existence: in each case strict equilibria are shown to exist.

We consider three different types of perturbations. First, based loosely on the earlier discussions of perturbations of beliefs and safety, we consider the possibility that group beliefs are random. Second, we consider the possibility that incentive constraints can be overcome by a costly enforcement technology. Finally, we suppose that group decisions are taken by a leader who has valence in the sense of being able to persuade group members to do as he wishes, but that if he issues orders that are not followed he is punished. In each case we take a limit: as beliefs become less random, enforcement becomes more costly, or valence shrinks; and in each case we show that the limit of equilibria of the perturbed games are collusion constrained equilibria in the unperturbed game. We emphasize that these are upper hemicontinuity results that do not show that every collusion constrained equilibrium arises this way. The issue of lower hemicontinuity is considered subsequently.

4.1. Random Belief Equilibrium

We now show that collusion constrained equilibria are limit points of strict collusion constrained equilibria when beliefs of each group about behavior of the other groups are

²⁴These types of equilibria as well as others where groups have more sophisticated correlating devices for internal use have been investigated in the context of voting models by Pogorelskiy (2014).

²⁵Despite the close relationship, the existence of collusion constrained equilibrium does not follow from the existence of quasi-equilibrium in Myerson (1982) nor can we use his argument since he assumes that principals have finitely many choices, while our groups choose from a continuum.

random and the randomness tends to vanish. We start by describing a random belief model. The idea is that given the true play ρ^{-k} of the other groups, there is a common belief σ^{-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 gives rise to common knowledge - looking each other in the eye, a handshake and so forth - 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 5. A density function $f^k(\sigma^{-k}|\rho^{-k})$ is called a *random group belief model* if it is continuous as a function of (σ^{-k}, ρ^{-k}) ; for $\epsilon > 0$ we say that the random group belief model is *only ϵ -wrong* if it satisfies $\int_{|\sigma^{-k} - \rho^{-k}| \leq \epsilon} f^k(\sigma^{-k}|\rho^{-k}) d\sigma^{-k} \geq 1 - \epsilon$.

In other words if the model is only ϵ -wrong then it places a low probability on being far from the truth. In Web Appendix 1 we give for every positive ϵ an example based on the Dirichlet distribution of a random group belief model that is only ϵ -wrong. We also define

Definition 6. A *group decision rule* is a function $b^k(\rho^{-k}) \in H[B^k(\rho^{-k})]$, measurable as a function of ρ^{-k} .

Notice that for given beliefs ρ^{-k} we are assuming that the group colludes on a response in $B^k(\rho^{-k})$ which is the set of the best choices for the group that satisfy the incentive constraints, and does not choose points in $B_S^k(\rho^{-k}) \setminus B^k(\rho^{-k})$ as would be permitted by shadow mixing.


Definition 7. For a group decision rule b^k and random group belief model f^k the *group response function* is the distribution $F^k(\rho^{-k})[a^k] = \int b^k(\sigma^{-k})[a^k] f^k(\sigma^{-k}|\rho^{-k}) d\sigma^{-k}$. If we have rules and belief models for all groups then a $\rho \in R$ that satisfies $\rho^k = F^k(\rho^{-k})$ for all k is called a *random belief equilibrium* with respect to b^k and f^k .

In the Appendix the following is proved:

Theorem 2. *Given for each k, n, ϵ_n group decision rules b^k and random group belief models $f_{\epsilon_n}^k$ that are only ϵ_n -wrong there are random belief equilibria ρ_n with respect to b^k and $f_{\epsilon_n}^k$. Moreover if $\epsilon_n \rightarrow 0$ and $\rho_n \rightarrow \rho$ then ρ is a collusion constrained equilibrium.*

Example (Random belief equilibrium in example 1). In Web Appendix 2 we analyze the random belief model corresponding to the Dirichlet belief model defined in Web Appendix

1. The figure below shows what the group response functions look like in our three player example. The key point is that the random belief equilibrium value of α^3 lies below $1/2$, that is, as $\epsilon \rightarrow 0$ the collusion constrained equilibrium is approached from the left and above.



4.2. Costly Enforcement Equilibrium

We now assume that each group k has a costly enforcement technology that it can use to overcome incentive constraints. In particular, we assume that every plan α^k is incentive compatible provided that the group pays a cost $C(\alpha^k, \rho^{-k})$ of carrying out the monitoring and punishment needed to prevent deviation. Levine and Modica (2016) show how cost of this type arise from peer discipline systems and Levine and Mattozzi (2016) study these systems in the context of voting by collusive parties: we give an example below. We assume $C^k(\alpha^k, \rho^{-k})$ to be non-negative and continuous in α^k, ρ^{-k} and adopt the following

Definition 8. A function $C^k(\alpha^k, \rho^{-k})$ is an *enforcement cost* if $C^k(\alpha^k, \rho^{-k}) = 0$ whenever $G^k(\alpha^k, \rho^{-k}) = 0$.

In other words enforcement is costly only if there is a deviation that needs to be deterred. Moreover, since nearby plans have similar gains to deviating, we assume also that the cost of deterring those deviations is similar - that is we assume that enforcement costs are continuous. . A particular example of such a cost function would be $C^k(\alpha^k, \rho^{-k}) = G^k(\alpha^k, \rho^{-k})$, that is, the cost of deterring a deviation is proportional to the biggest benefit any player receives by deviating. We give below an alternative example based on a technology for monitoring deviations. Notice that we allow the possibility that incentive incompatible plans have zero cost.

With this technology we define

Definition 9. The *enforced group best response set* $B_C^k(\rho^{-k})$ is the set of plans α^k such that $v^k(\alpha^k, \rho^{-k}) - C^k(\alpha^k, \rho^{-k}) = \max_{\tilde{\alpha}^k \in \mathcal{A}^k} v^k(\tilde{\alpha}^k, \rho^{-k}) - C^k(\tilde{\alpha}^k, \rho^{-k})$.

Notice that again there is no shadow mixing here, just a choice of the group's best plan. Then we have the usual definition of equilibrium

Definition 10. $\rho \in R$ is a *costly enforcement equilibrium* if $\rho^k \in H[B_C^k(\rho^{-k})]$.

Notice that if the cost of enforcement is zero then the group can achieve the best outcome ignoring incentive constraints, an assumption, as we indicated in the introduction, often used by political economists and economic historians. We are interested in the opposite case in which enforcing non-incentive compatible plans is very costly. We then define

Definition 11. A sequence $C_n^k(\alpha^k, \rho^{-k})$ of cost functions is *high cost* if there are sequences $\gamma_n^k \rightarrow 0$ and $\Gamma_n^k \rightarrow \infty$ such that $G^k(\alpha^k, \rho^{-k}) > \gamma_n^k$ implies $C_n^k(\alpha^k, \rho^{-k}) \geq \Gamma_n^k$.

In the Appendix we prove²⁶

Theorem 3. *Suppose $C_n^k(\alpha^k, \rho^{-k})$ is a high cost sequence. Then for each n a costly enforcement equilibrium ρ_n exists, and if $\lim_{n \rightarrow \infty} \rho_n \rightarrow \rho$ then ρ is a collusion constrained equilibrium.*

Example 2. We give a simple example of a costly enforcement technology and a high cost sequence based on Levine and Modica (2016). Specifically, we view the choice of α^k by group k as a social norm and assume that the group has a monitoring technology which generates a noisy signal of whether or not an individual member i complies with the norm. The signal is $z^i \in \{0, 1\}$ where 0 means “good, followed the social norm” and 1 means “bad, did not follow the social norm.” Suppose further that if member i violates the social norm by choosing $\alpha^i \neq \alpha^k$ then the signal is 1 for sure while if he adhered to the social norm so that $\alpha^i = \alpha^k$ then the signal is 1 with probability π_n . When the bad signal is received the group member receives a punishment of size P^i .²⁷

It is convenient to define the individual version of the gain to deviating

$$G^i(\alpha^k, \rho^{-k}) = \max_{d^i \in D^i} \sum_{a^{-k}} \left(U^i(d^i, \alpha^k, a^{-k}) - U^i(0, \alpha^k, a^{-k}) \right) \prod_{j \neq k} \rho^j[a^j] \geq 0.$$

For the social norm α^k to be incentive compatible we need $P^i - \pi_n P^i \geq G^i(\alpha^k, \rho^{-k})$ which is to say $P^i \geq G^i(\alpha^k, \rho^{-k}) / (1 - \pi_n)$. If the social norm is adhered to, the social cost of the punishment is $\pi_n P^i$, and the group will collude to minimize this cost so that it will choose $P^i = G^i(\alpha^k, \rho^{-k}) / (1 - \pi_n)$. The resulting cost is then $(\pi_n / (1 - \pi_n)) G^i(\alpha^k, \rho^{-k})$. Hence in this model $C_n^k(\alpha^k, \rho^{-k}) = (\pi_n / (1 - \pi_n)) \sum_{k(i)=k} G^i(\alpha^k, \rho^{-k})$.

²⁶Actually it is not essential that $\Gamma_n^k \rightarrow \infty$, just that it be “big enough” that it would never be worth paying such a high cost.

²⁷Here the coercion takes the form of punishment - but it could equally well be the withholding of a reward.

Since $C_n^k(\alpha^k, \rho^{-k}) = 0$ if and only if $G^k(\alpha^k, \rho^{-k}) = \max_{i|k(i)=k} G^i(\alpha^k, \rho^{-k}) = 0$ it follows that $C_n^k(\alpha^k, \rho^{-k})$ is an enforcement cost. We claim that as $\pi_n \rightarrow 1$, that is, as the signal quality deteriorates, this is in fact a high cost sequence. Certainly $C_n^k(\alpha^k, \rho^{-k}) \geq (\pi_n/(1 - \pi_n)) G^k(\alpha^k, \rho^{-k})$. Choose $\gamma_n^k \rightarrow 0$ such that $\Gamma_n^k \equiv (\pi_n/(1 - \pi_n)) \gamma_n^k \rightarrow \infty$. Then for $G^k(\alpha^k, \rho^{-k}) > \gamma_n^k$ we have $C_n^k(\alpha^k, \rho^{-k}) \geq \Gamma_n^k$ as required by the definition.

Example (Costly enforcement equilibrium in example 1). We use the high cost sequence just defined. In Web Appendix 2 we show that the costly enforcement equilibrium of our three-player game consists of the group randomizing half half between CC and DD while player 3 plays $\alpha^3 = (4 - 3\pi_n)/2$ for all $\pi_n > 4/5$. This equilibrium converges to the collusion constrained equilibrium as $\pi_n \rightarrow 1$. Notice that the collusion constrained equilibrium value of $\alpha^3 = 1/2$ is approached from the right while the group randomization in the costly enforcement equilibrium is constant and equal to the limiting constrained equilibrium value. This is the opposite of what we have seen in the random belief model where the approach is from the left and above.

4.3. Leader/Evaluator Equilibrium

In this section we tackle collusion constrained equilibrium from the perspective of the Nash program: showing how this partially cooperative notion arises from a limit of standard non-cooperative games. We do so by introducing leaders. Leaders give their followers instructions - they tell them things such as “let’s go on strike” or “let’s vote against that candidate.” The idea is that group leaders serve as explicit coordinating devices for groups. Each group will have a leader who tells group members what to do, and if he is to serve as an effective coordinating device these instructions cannot be optional. 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. Indeed, we do observe in practice that it is often the case that groups follow orders given by a leader but engage in *ex post* evaluation of the leader’s performance.

The leader/evaluator game is governed by two positive parameters ν, P . The parameter ν measures the “valence” of a leader: this has a concrete interpretation as the amount of utility that group members are ready to give up to follow the leader.²⁸ Alternatively, ν can be thought of as measuring group loyalty. The parameter P represents a punishment

²⁸It is convenient notationally and for the statement of results that all leaders have the same valence; this also implicitly assumes that followers of a leader are equally willing to sacrifice. This entails no loss of generality since as long as the willingness to sacrifice is positive we can linearly rescale u^i to units in which willingness to sacrifice is the same.

that can be levied by a group member against the leader.²⁹ Provided P is large enough, we will show that when valence tends to zero the limits of perfect Bayesian equilibria of the leader/evaluator game are collusion constrained equilibria of the original game.

Our non-cooperative game goes as follows:

Stage 1: Each leader chooses a plan $\alpha^k \in \mathcal{A}^k$ that is communicated only to members of group k : conceptually these are orders given to the members who must obey them.

Stage 2: Each player i with $k(i) = k$ serves as an evaluator and observing the plan α^k of the leader selects an element $d^i \in A^k \cup \{0\}$.

Payoffs: Let Q^k denote the number of evaluators who chose $d^i \neq 0$. The leader receives $v^k(\alpha^k, \alpha^{-k}) - PQ^k$, that is, for each evaluator who disagrees with his decision he is penalized by P . The evaluator receives utility $U^i(d^i, \alpha^k, \alpha^{-k})$ if $d^i \neq 0$ and $U^i(0, \alpha^k, \alpha^{-k}) + \nu$ if $d^i = 0$, that is, he takes as given the other players in the group have followed orders and gets a bonus of ν agrees with the leader's decision.

Note that the leader and evaluator do not learn what the other groups did until the game is over. In interpreting this game it is important to realize that the actions taken by group members are those ordered by the leader: the choice they make as evaluators are simply a statements of regret. So, for example, if the leader recommended a mixed strategy: mix 50-50 between C and D the choice $d^i = 0$ is a statement by the evaluator of satisfaction with that recommendation and the choice $d^i = C$ is a statement of dissatisfaction, the evaluator regrets not having chosen C . If the plan of the leader is regretted the evaluator then imposes a punishment on the leader.

Definition 12. We say that ρ is a *perfect Bayesian equilibrium* of the leader/evaluator game if for each leader k there is a mixed plan μ^k over \mathcal{A}^k , and for each evaluator i in each group k and each plan α^k there is a mixed action $\eta^i(\alpha^k)$ over $A^k \cup \{0\}$, measurable as a function of α^k , such that

- (i) $\rho^k = \int \sigma^k \mu^k(d\sigma^k)$
- (ii) μ^k (that is to say ρ^k) is optimal for the leader given ρ^{-k} and η^i
- (iii) for all $\alpha^k \in \mathcal{A}^k$ and evaluators i the measure $\eta^i(\alpha^k)$ is optimal for the evaluator given α^k and ρ^{-k} .

Note that (iii) embodies the idea of “no signaling what you do not know”³⁰ that beliefs about the play of leaders of other groups is independent of the plan chosen by the leader

²⁹Again this might depend on k but we can rescale ν^k so that punishment is the same for all leaders.

³⁰It is known for finite games that this is an implication of sequentiality and Fudenberg and Tirole (1991) use this condition to define perfect Bayesian equilibrium for a class of games. Since the leader/evaluator game is not finite sequentiality is complicated. Hence it seems most straightforward to follow Fudenberg and Tirole (1991) and define perfect Bayesian directly with the “no signaling what you do not know” condition.

of the own group.³¹ Note that we have not explicitly defined a system of beliefs, since the “no signaling what you do not know” condition makes the beliefs of evaluators over α^{-k} constant across all their information sets.

For this game to have an interesting relation to collusion constrained equilibrium, two things should be true.

- The evaluators must be able to punish the leader enough to prevent him from choosing incentive incompatible plans. A sufficient condition is that the punishment is greater than any possible gain in the game, that is, $P > \max v^k(\alpha^k, \alpha^{-k}) - \min v^k(\alpha^k, \alpha^{-k})$.
- The leader should be able to avoid punishment by choosing an incentive compatible plan. However the leader can only guarantee avoiding punishment if the evaluators strictly prefer not to deviate from his plan. If $\nu = 0$ this is true only for plans that are strictly incentive compatible and such plans may not exist. Hence the assumption $\nu > 0$ is crucial: it assures that the leader can always avoid punishment by choosing an incentive compatible plan.

The following result is proved in the Appendix.

Theorem 4. *Suppose $\nu_n \rightarrow 0$ and $P_n > \max v^k(\alpha^k, \alpha^{-k}) - \min v^k(\alpha^k, \alpha^{-k})$. Then for each n a perfect Bayesian equilibrium ρ_n of the leader/evaluator game exists, and if $\lim_{n \rightarrow \infty} \rho_n = \rho$ then ρ is a collusion constrained equilibrium.*

Example (Leader/Evaluator equilibrium in example 1). For $\alpha^3 < 1/2$, playing CC is incentive compatible for the group, the question is how much can they mix out of the unique bad within-group equilibrium DD when $\alpha^3 > 1/2$ given that they are willing to forgo gains no larger than ν . Web Appendix 2 shows that the equilibrium is $\hat{\alpha}^3 = (2 + \nu)/4 > 1/2$ and that the group mixes between the unique mixture $\hat{\alpha}^1 = \hat{\alpha}^2$ that is the smallest solution of $-4(\hat{\alpha}^1)^2(1 - \hat{\alpha}^3) + 2\hat{\alpha}^1 = \nu$ and CC with probability

$$\frac{0.5 - (\hat{\alpha}^1)^2}{1 - (\hat{\alpha}^1)^2}$$

on CC . Note that as $\nu \rightarrow 0$ we have $\hat{\alpha}^1 \rightarrow 0$ so that in the limit the group shadow mixes between CC and DD as expected. Notice also that $\alpha^3 > 1/2$ so that the solution is on “the same side” of $1/2$ as the costly enforcement equilibrium, but the opposite side of the belief equilibrium. The solution differs from both, however, in that the group does not randomize between CC and DD , but rather between CC and a mixed strategy.

³¹Since the leader has no way of knowing if other leaders have deviated he should not be able to signal this through his own choice of action.

5. Limits of Perturbations

In the perturbations we have considered the result is always that the limit of the perturbation is a collusion constrained equilibria. If there are several such equilibria, do the different limits converge to the same equilibrium? Not always. In this section we present an example with a continuum of collusion constrained equilibria and in which different perturbations pick different points out of this set.

The example is a variation of Example 1, where player 3 gets zero for sure if he plays C , and the good within-group equilibrium in the coordination game for the group which results if player 3 plays D is only weakly incentive compatible. We continue to set $v^1(a^1, a^2) = u^1(a^1, a^2) + u^2(a^1, a^2)$ and $v^2(a^1, a^2) = u^3(a^1, a^2)$.

Example 3. The matrix on the left below contains the payoffs if player 3 plays C , the right one results if she plays D :

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

In this game clearly player 3 must play D with probability 1: if he plays C with any positive probability then it is strictly dominant for players 1 and 2 to play D in which case player 3 strictly prefers to play D . When player 3 plays D players 1 and 2 have exactly two within-group equilibria: CC and DD ; and any mixture between them is a collusion constrained equilibrium. To see this observe that for any belief perturbation around $\alpha^3 = 0$ the worst within-group equilibrium for the group is always DD so $V_S^1(\alpha^3 = 0)$ is the utility the group obtains in that within-group equilibrium. Thus any mixture between DD and CC satisfies the equilibrium condition, where of course in all strictly mixed equilibria the group gets utility higher than V_S^1 .

Now consider the perturbations. For any random beliefs C has positive probability so the group must play DD , so the only limit of random belief equilibria is DD . For costly enforcement equilibrium on the other hand the better within-group equilibrium CC for the group has zero cost so that will be chosen: the unique limit in this case is CC . Finally, for leadership equilibrium since the compliance bonus ν is positive again CC will be chosen, the unique limit is again CC . Notice that not only do the different perturbations sometimes pick different points out of the collusion constrained equilibrium set, but the collusion constrained equilibria involving strict mixtures do not arise as a limit from any of the perturbed models.

This example is non-generic because it depends heavily on the fact that when player 3 plays a pure strategy D players 1 and 2 are indifferent to deviating from CC . If we try to construct an example of this type in the interior then players 1 and 2 must shadow mix

in the correct way to make player 3 indifferent and this should pin down what the shadow mixture must be. In the example we get around this by assuming that the pure strategy for player 3 is a strict best response so that there are a continuum of shadow mixtures by 1 and 2 that are consistent with player 3 playing D .³²

6. Lower Hemicontinuity and a Characterization of CCE

Roughly speaking, when we consider a perturbation such as random belief equilibrium, leadership equilibrium, or costly enforcement equilibrium we are exhibiting a degree of agnosticism about the model we have written down. That is we recognize that our model is an imperfect representation but hopefully reasonable approximation of a more complex reality and ask whether our equilibrium might be a good description of what happens in that more complex reality. This is the spirit behind refinements such as trembling hand perfection and concepts such as Harsanyi (1973)'s notion of purification of a mixed equilibrium. It is the question addressed by Fudenberg, Kreps and Levine (1988) who show how refinements do not capture the equilibria of all nearby games. We have shown that collusion constrained equilibrium does a good job of capturing random beliefs, costly enforcement and leadership equilibria. We know by example that there may be collusion constrained equilibria that do not arise as a limit of any of these. We now ask whether for a given collusion constrained equilibrium there is a story we can tell in the form of a perturbation representing a more complex reality that justifies the particular collusion constrained equilibrium.

Each of the perturbations we have considered has embodied a story or justification about why groups might be playing the way they are playing. We now consider a richer class of perturbations that combines elements of beliefs with costly enforcement and a perturbation of the group objective function. Specifically, we use the following:

Definition 13. A *perturbation* for each group k consists of a continuous belief perturbation $r_k^{-k}(\rho^{-k}) \in R^{-k}$, an enforcement cost function $C^k(\alpha^k, \rho^{-k})$ and a continuous objective function $w^k(\alpha^k, \rho^{-k})$. A *perturbed equilibrium* ρ is defined by the condition $\rho^k \in H[\arg \max_{\alpha^k} w^k(\alpha^k, r_k^{-k}(\rho^{-k})) - C^k(\alpha^k, r_k^{-k}(\rho^{-k}))]$.

The belief perturbation is a simplification of the random belief model which assumed that beliefs were random but near correct most of the time. Now we are going to assume that they are deterministic and near correct. As in the random belief model we allow that beliefs are slightly wrong and do not require that two groups agree about the play of a

³²We do not know if generic examples exist - genericity is quite difficult to analyze in this model. That our results on lower hemi-continuity in the next section make use of perturbations of the group utility function suggests that examples in which the limits fail to coincide may well be non-generic.

third. The model of costly enforcement is exactly the same model we studied earlier. In addition we are now agnostic about the group objective and allow the possibility that the model may be slightly wrong in this respect. From a technical point of view it helps get rid of non-generic examples. As we are only interested in small perturbations we define

Definition 14. A sequence of perturbations $r_{kn}^{-k}, C_n^k, w_n^k$ is said to *converge* as $n \rightarrow \infty$ if $\max_{\rho^{-k}} |r_{kn}^{-k}(\rho^{-k}) - \rho^{-k}| \rightarrow 0$, if C_n^k is a high cost sequence, and if $\max_{\alpha^k, \rho^{-k}} |w_n^k(\alpha^k, \rho^{-k}) - v^k(\alpha^k, \rho^{-k})| \rightarrow 0$. We say that ρ is *justifiable* if there is a convergent sequence of perturbations together with perturbed equilibria $\rho_n \rightarrow \rho$.

Our main result, proven in the Appendix, is

Theorem 5. *A perturbed equilibrium exists for any perturbation, and ρ is justifiable if and only if it is a collusion constrained equilibrium.*

7. A Voting Participation Game

What difference do groups make? Collusion constrained equilibria are a subset of the set of group correlated equilibria, so we should expect that often the equilibria that are rejected are going to have better efficiency properties than those that are accepted. However, that comparison is not so interesting because it is the fact that the group is collusive that enables it to randomize privately from the other groups - that is, coordinate their play.³³ A more useful comparison is to ask what happens if the players play as individuals without correlating devices to coordinate their play, versus what happens if they are in collusive groups. So, in addition to static Nash equilibrium a second useful benchmark is to analyze the case in which there is free (costless) enforcement (FEE) - so that incentive constraints do not matter.

Our setting for studying the economics of collusion is a voter participation game. We start with a relatively standard Palfrey and Rosenthal (1985) framework (see also Levine and Mattozzi (2016)): there are two parties, the “large” party has two voters, players 1 and 2, the “small” party has one voter, player 3. Voters always vote for their own party, but it is costly to vote - a cost we normalize to 1 - and voters may choose whether or not to vote. The party that wins receives a transfer payment of $2\tau > 0$ from the losing party: if the large party wins player 3 loses 2τ which the large party members split; if player 3 wins she gets τ from each member of the large party. Usually it is assumed that a tie means that each party has a 50% chance of winning the prize, meaning that the election is a wash and

³³The random belief model, in particular, only makes sense if the group is colluding, otherwise how can they agree on their beliefs?

no transfer payment is made. In case nobody votes we maintain this assumption that the status quo is unchanged and everyone gets 0. But when voting does take place it is often not the case in practice that a tie is innocuous - it may result in a deadlocked government or in conflict between the parties. So we assume that a tie where each party casts one vote results in a deadlock that is - for simplicity - just as bad as a loss. The group objective of either party is to maximize the sum of payoffs of its members.

The payoffs can be written in bi-matrix form. If player 3 does not vote the payoff matrix for the actions of players 1 and 2 (where 0 represents do not vote and 1 represents vote) is

$$\begin{array}{cc}
 & \begin{array}{c} 1 \\ 0 \end{array} & \begin{array}{c} 0 \\ 0, 0, 0 \end{array} \\
 \begin{array}{c} 1 \\ 0 \end{array} & \begin{array}{c} \tau - 1, \tau - 1, -2\tau \\ \tau, \tau - 1, -2\tau \end{array} & \begin{array}{c} \tau - 1, \tau, -2\tau \\ 0, 0, 0 \end{array}
 \end{array}$$

This game between players 1 and 2 has a unique dominant strategy equilibrium at which neither votes if $\tau < 1$. If player 3 does vote the payoff matrix for the actions of players 1 and 2 becomes

$$\begin{array}{cc}
 & \begin{array}{c} 1 \\ 0 \end{array} & \begin{array}{c} 0 \\ -\tau, -\tau, 2\tau - 1 \end{array} \\
 \begin{array}{c} 1 \\ 0 \end{array} & \begin{array}{c} \tau - 1, \tau - 1, -2\tau - 1 \\ -\tau, -\tau - 1, -2\tau - 1 \end{array} & \begin{array}{c} -\tau - 1, -\tau, -2\tau - 1 \\ -\tau, -\tau, 2\tau - 1 \end{array}
 \end{array}$$

If $\tau > 1/2$ this is a coordination game for player 1 and 2 due to the fact that a tie is as bad as a loss: for a large party member not voting and having a tie results in $-\tau$ while voting and winning results in $\tau - 1 > -\tau$. Similarly voting and having a tie is as bad as a loss and it would be better to not vote and lose, suffering the same loss but not paying the cost of voting.

The model has elements of both external and internal conflict. There is conflict between the groups as each hopes to get the transfer. There is also conflict within the large group as each prefers that the other votes. There are two sources of inefficiency in the model: total welfare (the sum of the utilities of the all three players) is reduced if players vote and is further reduced if there is a tie with one player from each group voting.

The full analysis of the structure of collusion constrained, Nash and free enforcement equilibria in this model can be found in Web Appendix 3. To appreciate the usefulness of CCE, focus on the range $3/4 < \tau < 1$. Here there is a unique Nash equilibrium S in which only the small group votes and a unique FEE L in which the small group abstains and the large group wins by casting a single vote. In this range there is also a unique CCE in which the small group mixes on voting and not voting with positive probability and the large group shadow mixes between staying out with probability $1/2\tau$ and casting two votes. In the CCE equilibrium the small group does better than at FEE and worse than Nash while

the large group does worse than at FEE and better than Nash. The CCE more accurately captures the behavior of a collusive group as one that is in between the Nash prediction of extreme free riding and the FEE prediction of complete disregard of individual incentives. A more subtle implication relates to the equilibrium behavior of the small party. Despite consisting of a single player, the CCE aptly captures how equilibrium behavior depends significantly on whether the player faces an individual or a group, and in the latter case then whether it is collusive or not.

Varying τ provides a richer but similar picture. First note that among all equilibria of all types, when they are equilibria S is always best for the small group and L for the large group. Start with $\tau < 1/2$ in which case nobody votes. As we increase τ Nash always allows S although for $\tau > 1$ there are additional equilibria less favorable to the small player, including L . CCE and FEE both shift gradually in favor of the large group but CCE changes more slowly than does FEE: for FEE once $\tau > 3/4$ the unique equilibrium is L while for CCE this is true only for $\tau > 3/2$.

8. Conclusion

We study exogenously specified collusive groups and argue that the “right” notion of equilibrium is that of collusion constrained equilibrium. We start from the observation that groups such as political, ethnic, business or religious groups often collude. We adopt the simple assumption that a group will collude on the within-group equilibrium that best satisfies group objectives. 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 “fix” 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 within-group equilibrium the incentive constraints are satisfied with equality. This “tremble” implies that the group may put positive probability on actions which are worse for the group but are strictly incentive compatible. We show that the resulting equilibrium notion has strong robustness properties and indeed is both upper and lower hemicontinuous with respect to a class of perturbations. This makes collusion constrained equilibrium a strong foundation 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 Olson (1965) had in mind and is of key importance in much of the political economy literature. This is not to argue that endogenous group formation is not of interest - but it is important to understand what happens as a consequence of group formation before building models of group formation and collusion constrained equilibrium is a step in that direction.

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.
- Aliprantis, Charalambos and Kim Border (2007): *Infinite Dimensional Analysis: A Hitchhiker's Guide (3rd Edition)*, Springer
- Arrow, K. J. and G. Debreu (1954): "Existence of an equilibrium for a competitive economy," *Econometrica* 22: 265-290.
- Balasubramanian, S., and P. Bhardwaj (2004): "When not all conflict is bad: Manufacturing-marketing conflict and strategic incentive design," *Management Science* 50: 489-502.
- Bernheim, B.Douglas, Bezalel Peleg, Michael D Whinston (1987): "Coalition-Proof Nash Equilibria I. Concepts," *Journal of Economic Theory* 42: 1-12.
- Bich, Philippe and Rida Laraki (2016): "On the Existence of Approximate Equilibria and Sharing Rule Solutions in Discontinuous Games," *Theoretical Economics*, forthcoming.
- Bolton, Patrick, Markus Brunnermeier and Laura Veldkamp (2013): "Leadership, Coordination and Corporate Culture", *Review of Economic Studies* 80: 512-537
- Caillaud, Bernard and Philippe Jehiel (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* 44: 69-95.
- 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* 144: 565-603.
- 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.
- Fershtman, Chaim, and Kenneth L. Judd (1986): "Strategic Incentive in Manipulation in Rivalrous Agency," Institute for Mathematical Studies in the Social Sciences, Stanford University.

- Fudenberg, Drew, David Kreps and David Levine (1988): "On the robustness of Equilibrium Refinements", *Journal of Economic Theory* 44: 354-380.
- Fudenberg, Drew and Jean Tirole (1991): Perfect Bayesian equilibrium and sequential equilibrium, *Journal of Economic Theory* 53: 236-260
- Haeringer, Guillaume. (2004): "Equilibrium Binding Agreements: A Comment," *Journal of Economic Theory* 117: 140-143.
- Harrington, Joseph. (2008): "Detecting Cartels," in P. Buccirosi, ed., *Handbook of Antitrust Economics*, MIT Press.
- 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.
- Kopel, Michael, and Clemens Löffler (2012): "Organizational governance, leadership, and the influence of competition," *Journal of Institutional and Theoretical Economics* 168: 362-392.
- T. Hu, J. Kennan and N. Wallace (2009): "Coalition-Proof Trade and the Friedman Rule in the Lagos-Wright Model," *Journal of Political Economy* 117: 116-137.
- Laffont, Jean-Jacques and David Martimort (1997): "Collusion Under Asymmetric Information," *Econometrica* 65: 875-911.
- Levine, David and Salvatore Modica (2016): "Peer Discipline and Incentives within Groups", *Journal of Economic Behavior and Organization*, in press.
- Levine, David and Andrea Mattozzi (2016): "Voter Participation with Collusive Parties," <http://www.dklevine.com/archive/refs478696900000001234.pdf>.
- Martimort, D. and H. Moreira (2010): "Common Agency and Public Good Provision under Asymmetric Information", *Theoretical Economics* 5: 159-213.
- 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.
- Myerson, Roger B. (1982): "Optimal coordination mechanisms in generalized principal-agent problems," *Journal of Mathematical Economics* 10: 67-81.
- Olson, Mancur (1965): *The Logic of Collective Action: Public Goods and the Theory of Groups*.
- Palfrey, Thomas R., and Howard Rosenthal (1985): "Voter participation and strategic uncertainty," *American Political Science Review* 79: 62-78.

Pogorelskiy, K. (2014): “Correlated Equilibria in Voter Turnout Games,” Warwick.

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

Ray, Debraj and Rajiv Vohra (1997): “Equilibrium Binding Agreements,” *Journal of Economic Theory* 73: 30-78.

Reny, Philip J. (1999): “On the Existence of Pure and Mixed Strategy Nash Equilibria in Discontinuous Games,” *Econometrica* 67, 1029-1056.

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: Continuity, Limits and Existence

Lemma 1. *Suppose we have a sequence of sets B_n^k , correlated profiles $\rho_n^{-k} \rightarrow \rho^{-k}$, scalars V_n^k and positive numbers $\gamma_n^k \rightarrow 0$ satisfying for any $\alpha_n^k \in B_n^k$*

1. $G^k(\alpha_n^k, \rho_n^{-k}) \leq \gamma_n^k$

2. $v^k(\alpha_n^k, \rho_n^{-k}) \geq V_n^k$

If B^k is the set of $\alpha^k \in B^k$ that satisfies

1. $G^k(\alpha^k, \rho^{-k}) = 0$

2. $v^k(\alpha^k, \rho^{-k}) \geq \liminf V_n^k$

then for any $\rho_n^k \in H(B_n^k)$ with $\rho_n^k \rightarrow \rho^k$ it is the case that $\rho^k \in H(B^k)$

Proof. Since G^k, v^k are continuous the closure of B_n^k satisfies the same inequalities so it suffices to prove the result for closed sets B_n^k .

We have $\rho_n^k \in H(B_n^k)$ if and only if there exists a probability measure μ_n^k over B_n^k with $\rho_n^k = \int \sigma \mu_n^k(d\sigma)$. Since B_n^k is closed $\mathcal{A}^k \setminus B_n^k$ is open and we can extend the measure to all of \mathcal{A}^k by taking $\mu_n^k[\mathcal{A}^k \setminus B_n^k] = 0$. Since \mathcal{A}^k is compact we may extract a weakly convergent subsequence that converges to μ^k and without loss of generality may assume the original sequence has this property. Because $\mu_n^k \rightarrow \mu^k$ it follows from weak convergence that $\rho^k = \int \sigma \mu^k(d\sigma)$. The result will follow if we can show that $\mu^k[B^k] = 1$.

Consider the sets B_v^k for which $v^k(\alpha^k, \rho^{-k}) \geq \liminf V_n^k$ and B_0^k for which $G^k(\alpha^k, \rho^{-k}) = 0$. We will show that $\mu^k[B_v^k] = 1$ and $\mu^k[B_0^k] = 1$ from which it follows that $\mu^k[B^k] = \mu^k[B_v^k \cap B_0^k] = 1$.

For B_v^k let $\epsilon > 0$ and let $D_{v\epsilon}^k$ be the set $v^k(\alpha^k, \rho^{-k}) < \liminf V_n^k - \epsilon$. For n sufficiently large $D_{v\epsilon}^k \cap B_n^k = \emptyset$, so $\mu_n^k[D_{v\epsilon}^k] = 0$. However since v^k is continuous $D_{v\epsilon}^k$ is an open set and

if $\mu^k[D_{v\epsilon}^k] > 0$ then for all sufficiently large n we have $\mu_n^k[D_{v\epsilon}^k] > 0$, a contradiction. We conclude that for all $\epsilon > 0$ we have $\mu^k[D_{v\epsilon}^k] = 0$, so indeed $\mu^k[B_v^k] = 1$.

For B_0^k let $\epsilon > 0$ and let $D_{0\epsilon}^k$ be the set $G^k(\alpha^k, \rho^{-k}) > \epsilon$. Because $\mathcal{A}^k \times R^{-k}$ is compact $G^k(\alpha^k, \rho^{-k})$ is uniformly continuous so $G^k(\cdot, \rho_n^{-k})$ converges uniformly to $G^k(\cdot, \rho^{-k})$. Hence for n sufficiently large $\alpha^k \in D_{0\epsilon}^k$ implies $G^k(\alpha^k, \rho_n^{-k}) > \epsilon/2$ and since $\gamma_n^k \rightarrow 0$ also for sufficiently large n this implies $\mu_n^k[D_{0\epsilon}^k] = 0$. However, since G^k is continuous $D_{0\epsilon}^k$ is an open set, and if $\mu^k[D_{0\epsilon}^k] > 0$ then for all sufficiently large n we have $\mu_n^k[D_{0\epsilon}^k] > 0$ a contradiction. We conclude that for all $\epsilon > 0$ we have $\mu^k[D_{0\epsilon}^k] = 0$, so indeed $\mu^k[B_0^k] = 1$. \square

Corollary 1. *Let the sets B_n^k be satisfy $G^k(\alpha_n^k, \rho_n^{-k}) \leq \gamma_n^k$ and $v^k(\alpha_n^k, \rho_n^{-k}) \geq V_{\epsilon_n}^k(\rho_n^{-k})$. If $\gamma_n^k, \epsilon_n \rightarrow 0$ and $\rho_n^k \in H(B_n^k) \rightarrow \rho^k$ for all k then ρ is a collusion constrained equilibrium.*

Proof. If $\epsilon_n \leq \epsilon/2$ and $|\rho_n^{-k} - \rho^{-k}| \leq \epsilon/2$ then $|\sigma_n^{-k} - \rho_n^{-k}| \leq \epsilon_n$ implies $|\sigma_n^{-k} - \rho^{-k}| \leq \epsilon$ whence $V_{\epsilon_n}^k(\rho_n^{-k}) \geq V_{\epsilon}^k(\rho^{-k})$. This gives $\liminf V_{\epsilon_n}^k(\rho_n^{-k}) \geq V_S^k(\rho^{-k})$. Therefore taking $V_n^k = V_{\epsilon_n}^k(\rho_n^{-k})$, Lemma 1 shows that ρ^k is contained in the convex hull of a set contained in $B_S^k(\rho^{-k})$ for all k , whence the conclusion. \square

Collusion Constrained Equilibrium

Theorem 6 (Theorem 1 in text). *Collusion constrained equilibria exist and are a subset of the set of group correlated equilibria of the game.*

Proof. For any sequence of correlated profiles $\rho_n^{-k} \rightarrow \rho^{-k}$, let $\gamma_n^k = 0$ and let $V_n^k = V_S^k(\rho_n^k)$. Notice that $\liminf V_n^k \geq V_S^k(\rho^k)$. Then by Lemma 1 we know that the convex hull of the shadow response set, $H(B_S^k(\rho^{-k}))$ is UHC. Existence of collusion constrained equilibria then follows from Kakutani. The fact that collusion constrained equilibria are group correlated equilibria follows from the fact that the incentive constraints are satisfied for each individual given signals generated by the private and group randomizing devices. \square

Random Belief Equilibria

Theorem 7 (Theorem 2 in text). *Given for each k, n, ϵ_n , group decision rules b^k and random group belief models $f_{\epsilon_n}^k$ that are only ϵ_n -wrong there are random belief equilibria ρ_n with respect to b^k and $f_{\epsilon_n}^k$. Moreover if $\epsilon_n \rightarrow 0$ and $\rho_n \rightarrow \rho$ then ρ is a collusion constrained equilibrium.*

Proof. Remember that $\rho_n^k(a^k) = F^k(\rho^{-k})[a^k] = \int b^k(\sigma^{-k})[a^k] f_{\epsilon_n}^k(\sigma^{-k}|\rho^{-k}) d\sigma^{-k}$ where $f_{\epsilon_n}^k(\sigma^{-k}|\rho^{-k})$ is continuous as a function of ρ^{-k} . So $\rho_n^k(a^k)$ is a continuous function of ρ^{-k} by the Dominated Convergence Theorem, for every a^k . Existence then follows from the Brouwer fixed point theorem.

Turning to convergence, by definition

$$\begin{aligned}\rho_n^k &= \int b^k(\sigma^{-k}) f_{\epsilon_n}^k(\sigma^{-k}|\rho^{-k}) d\sigma^{-k} \\ &= \int_{|\sigma^{-k}-\rho^{-k}|\leq\epsilon_n} b^k(\sigma^{-k}) f_{\epsilon_n}^k(\sigma^{-k}|\rho^{-k}) d\sigma^{-k} + \int_{|\sigma^{-k}-\rho^{-k}|>\epsilon_n} b^k(\sigma^{-k}) f_{\epsilon_n}^k(\sigma^{-k}|\rho^{-k}) d\sigma^{-k}\end{aligned}$$

Let $e_n^k(\rho^{-k}) \equiv \int_{|\sigma^{-k}-\rho^{-k}|\leq\epsilon_n} f_{\epsilon_n}^k(\sigma^{-k}|\rho^{-k}) d\sigma^{-k}$ and

$$\bar{\rho}_n^k \equiv \int_{|\sigma^{-k}-\rho^{-k}|\leq\epsilon_n} b^k(\sigma^{-k}) \frac{f_{\epsilon_n}^k(\sigma^{-k}|\rho^{-k})}{e_n^k(\rho^{-k})} d\sigma^{-k}$$

then we may write

$$\rho_n^k = e_n^k(\rho^{-k}) \bar{\rho}_n^k + (1 - e_n^k(\rho^{-k})) \int_{|\sigma^{-k}-\rho^{-k}|>\epsilon_n} b^k(\sigma^{-k}) \frac{f_{\epsilon_n}^k(\sigma^{-k}|\rho^{-k})}{1 - e_n^k(\rho^{-k})} d\sigma^{-k}.$$

Now assume $\epsilon_n \rightarrow 0$. By assumption $e_n^k(\rho^{-k}) \rightarrow 1$ and $\rho_n^k \rightarrow \rho^k$ it follows that $\bar{\rho}_n^k \rightarrow \rho^k$. Take then $B_n^k \equiv \{\alpha^k \in B^k(\sigma^{-k}) \mid |\sigma^{-k} - \rho_n^{-k}| \leq \epsilon_n\}$. Clearly $\bar{\rho}_n^k \in H(B_n^k)$. We now show that the sequence $(\bar{\rho}_n^k, \rho_n^{-k})$ satisfies the hypotheses of Corollary 1. For any $\alpha_n^k \in B_n^k$ there is σ_n^{-k} with $|\sigma_n^{-k} - \rho_n^{-k}| \leq \epsilon_n$ such that $G^k(\alpha_n^k, \sigma_n^{-k}) = 0$ and $v^k(\alpha_n^k, \sigma_n^{-k}) = V^k(\sigma_n^{-k})$. Taking

$$\gamma_n^k = \max_{\alpha^k \in \mathcal{A}^k} \max_{|\sigma^{-k}-\rho^{-k}|\leq\epsilon_n} |G^k(\alpha^k, \sigma^{-k}) - G^k(\alpha^k, \rho^{-k})|$$

we see that $G^k(\alpha_n^k, \rho_n^{-k}) \leq \gamma_n^k$. Since G^k is continuous on a compact set it is uniformly continuous so $\gamma_n^k \rightarrow 0$. Moreover if $\alpha_n^k \in B_n^k$ then clearly $v^k(\alpha_n^k, \rho_n^{-k}) \geq V_{\epsilon_n}^k(\rho_n^{-k})$. The result now follows from Corollary 1. \square

Leadership Equilibrium

For $\nu > 0$ define $V_\nu^k(\rho^{-k}) = \sup_{\alpha^k \in \mathcal{A}^k \mid G^k(\alpha^k, \rho^{-k}) < \nu} v^k(\alpha^k, \rho^{-k})$ and $B_\nu^k(\rho^{-k})$ to be the set of plans α^k satisfying $G^k(\alpha^k, \rho^{-k}) \leq \nu$ and $v^k(\alpha^k, \rho^{-k}) \geq V_\nu^k(\rho^{-k})$.

Definition 15. We say that ρ is a *strict ν -equilibrium* if $\rho^k \in H[B_\nu^k(\rho^{-k})]$ for all k .

Theorem 8. *Strict ν -equilibria exist.*

Proof. It is sufficient to show that B_ν^k is UHC. By Theorem 17.35 in Aliprantis and Border (2007) we then know that $H[B_\nu^k(\rho^{-k})]$ is also UHC. Existence of strict ν -equilibrium then follows by Kakutani's fixed point theorem.

Consider a sequence $(\alpha_n^k, \rho_n^{-k})$ such that $\alpha_n^k \in B_\nu^k(\rho_n^{-k})$. Suppose that $\lim_{n \rightarrow \infty} \alpha_n^k = \alpha^k$ and $\lim_{n \rightarrow \infty} \rho_n^{-k} = \rho^{-k}$. By continuity, $G^k(\alpha_n^k, \rho_n^{-k}) \leq \nu$ for all n implies that $G^k(\alpha^k, \rho^{-k}) \leq$

ν . Suppose by contradiction, $v^k(\alpha^k, \rho^{-k}) < V_\nu^k(\rho^{-k})$. By the continuity of v^k it follows that for sufficiently large n we have $v^k(\alpha_n^k, \rho_n^{-k}) < V_\nu^k(\rho^{-k})$. Since $v^k(\alpha_n^k, \rho_n^{-k}) \geq V_\nu^k(\rho_n^{-k})$ this implies $V_\nu^k(\rho_n^{-k}) < V_\nu^k(\rho^{-k})$. Hence there is some $\hat{\alpha}^k$ such that $G^k(\hat{\alpha}^k, \rho^{-k}) < \nu$ and $V_\nu^k(\rho_n^{-k}) < v^k(\hat{\alpha}^k, \rho^{-k})$. By continuity of G^k and v^k this in turn implies that for sufficiently large n we have $G^k(\hat{\alpha}^k, \rho_n^{-k}) < \nu$ and $V_\nu^k(\rho_n^{-k}) < v^k(\hat{\alpha}^k, \rho_n^{-k})$ contradicting the definition of $V_\nu^k(\rho_n^{-k})$. \square

Theorem 9. ρ is a perfect Bayesian equilibrium of the leader evaluator game if and only if it is a strict ν -equilibrium.

Proof. Suppose ρ is perfect Bayesian. Let μ^k and η^i be the corresponding leader and evaluator strategies. It suffices to show that $\mu^k[B_\nu^k(\rho^{-k})] = 1$. Denote the equilibrium utility of leader k by U^k .

Let D_ν^k be the subset of \mathcal{A}^k for which $G^k(\alpha^k, \rho^{-k}) > \nu$. For $\alpha^k \in D_\nu^k$ there is an i with $k(i) = k$ for whom it is optimal to choose $\eta^i(\alpha^k)[\alpha^k] = 0$, hence utility for the leader is at most $\max v^k - P$ for those choices of α^k . Suppose $d = \mu^k[D_\nu^k] > 0$. Let $\hat{\alpha}^k \in \mathcal{A}^k$ satisfy $G^k(\hat{\alpha}^k, \rho^{-k}) = 0$ which we know exists. Consider $\hat{\mu}^k$ that takes the weight from D_ν^k and puts it on $\hat{\alpha}^k$. The utility from $\hat{\mu}^k$ is at least $(1-d)U^k + d(U^k + \min v^k - \max v^k + P)$ which is bigger than U^k since $P > \max v^k - \min v^k$. Hence $d = 0$.

Let \tilde{D}_ν^k be the subset of \mathcal{A}^k for which $v^k(\alpha^k, \rho^{-k}) < V_\nu^k(\rho^{-k}) - \epsilon$. Suppose $\tilde{d} = \mu^k[\tilde{D}_\nu^k] > 0$. Let $\tilde{\alpha}^k \in \mathcal{A}^k$ satisfy $G^k(\tilde{\alpha}^k, \rho^{-k}) < \nu$ and $v^k(\tilde{\alpha}^k, \rho^{-k}) > V_\nu^k(\rho^{-k}) - \epsilon/2$ which we know exists. By evaluator optimality we have $\eta^i(\tilde{\alpha}^k)[\tilde{\alpha}^k] = 1$ for all $k(i) = k$. Consider $\tilde{\mu}^k$ that takes the weight from \tilde{D}_ν^k and puts it on $\tilde{\alpha}^k$. The utility from $\tilde{\mu}^k$ is at least $U^k + d\epsilon/2$ so $\tilde{d} = 0$. Since $B_\nu^k(\rho^{-k}) \subseteq D_\nu^k \cup \tilde{D}_\nu^k$ we see that that indeed $\mu^k[B_\nu^k(\rho^{-k})] = 1$.

Now suppose that ρ is a strict ν -equilibrium. Since $\rho^k \in H[B_\nu^k(\rho^{-k})]$ there exist measures μ^k with $\mu^k[B_\nu^k(\rho^{-k})] = 1$ and $\rho^k = \int \sigma \mu^k(d\sigma)$ so it suffices to find η^i that together with μ^k form a perfect Bayesian equilibrium. Specifically for the given ρ^k we show how to choose evaluator optimal responses η^i to each α^k such that the given μ^k are optimal for the leader with respect to those evaluations. Start by choosing an optimal default for the evaluators - a function $\bar{\alpha}^i(\alpha^k) \in \arg \max_{\alpha^i} u^i(\alpha^i, \alpha^k, \rho^{-k})$ that is measurable. We define evaluator optimal responses η^i to α^k by the leader in three cases depending on the size $G^k(\alpha^k, \rho^k)$ and $v^k(\alpha^k, \rho^{-k})$. Observe that it cannot be that $G^k(\alpha^k, \rho^{-k}) < \nu$ and $v^k(\alpha^k, \rho^{-k}) > V_\nu^k(\rho^{-k})$, so we may omit consideration of that case.

- (i) if $G^k(\alpha^k, \rho^{-k}) > \nu$ then $\eta^i[\bar{\alpha}^i(\alpha^k)] = 1$. Note that in this case $\bar{\alpha}^i(\alpha^k) \neq \alpha^i$ for at least one i
- (ii) if $G^k(\alpha^k, \rho^{-k}) \leq \nu$ and $v^k(\alpha^k, \rho^{-k}) \leq V_\nu^k(\rho^{-k})$ then $\eta^i[\alpha^i] = 1$
- (iii) if $G^k(\alpha^k, \rho^{-k}) = \nu$ and $v^k(\alpha^k, \rho^{-k}) > V_\nu^k(\rho^{-k})$ some evaluator j is indifferent between α^j and some $\bar{\alpha}^j \neq \alpha^j$ (and this evaluator can be chosen in a measurable way).

For $i \neq j$ take $\eta^i[\alpha^i] = 1$. For j choose $\eta^j[\bar{\alpha}^j] = (v^k(\alpha^k, \rho^{-k}) - V_\nu^k(\rho^{-k})) / P$ and $\eta^j[\alpha^j] = 1 - \eta^j[\bar{\alpha}^j]$.

We now establish that for this evaluator optimal response μ^k is indeed optimal for the leader. If $\alpha^k \in B_\nu^k(\rho^{-k})$ the leader utility is exactly $V_\nu^k(\rho^{-k})$, while if $G^k(\alpha^k, \rho^{-k}) > \nu$ then leader utility is at most $\max v^k - P$. Hence $\alpha^k \in B_\nu^k(\rho^{-k})$ is at least as good as any other choice, and indifferent to any other choice in $B_\nu^k(\rho^{-k})$. Since μ^k is a randomization over $B_\nu^k(\rho^{-k})$ for the leader it follows that it is optimal. \square

Lemma 2. $V_\nu^k(\rho^{-k}) \geq V_\epsilon^k(\rho^{-k})$ for any $\epsilon > 0$.

Proof. From

$$\begin{aligned} V_\nu^k(\rho^{-k}) &= \sup_{\alpha^k \in \mathcal{A}^k \mid G^k(\alpha^k, \rho^{-k}) < \nu} v^k(\alpha^k, \rho^{-k}) \\ &\geq \sup_{\alpha^k \in \mathcal{A}^k \mid G^k(\alpha^k, \rho^{-k}) = 0} v^k(\alpha^k, \rho^{-k}) = V^k(\rho^{-k}) \geq V_\epsilon^k(\rho^{-k}) \end{aligned}$$

the stated inequality follows. \square

Theorem 10. *If ρ_n is a sequence of strict ν_n -equilibria, $\nu_n \rightarrow 0$ and $\rho_n \rightarrow \rho$ then ρ is a collusion constrained equilibrium.*

Proof. Let $\gamma_n = \nu_n$ and notice that for any $\alpha_n^k \in B_{\nu_n}^k(\rho_n^{-k})$ we have $v^k(\alpha_n^k, \rho_n^{-k}) \geq V_{\nu_n}^k(\rho_n^{-k}) \geq V_{\epsilon_n}^k(\rho_n^{-k})$ by Lemma 2 for some sequence $\epsilon_n \rightarrow 0$. Result now follows from Corollary 1. \square

Perturbed Equilibrium: Existence and Upper HemiContinuity

Theorem 11. *A perturbed equilibrium exists for any perturbation.*

Proof. Notice that for any perturbation $w^k(\alpha^k, r_k^{-k}(\rho^{-k})) - C^k(\alpha^k, r_k^{-k}(\rho^{-k}))$ is continuous in its arguments. By the Maximum Theorem we then get the correspondence $\arg \max_{\alpha^k} w^k(\alpha^k, r_k^{-k}(\rho^{-k})) - C^k(\alpha^k, r_k^{-k}(\rho^{-k}))$ to be UHC. In turn by Theorem 17.35 in Aliprantis and Border (2007), $H[\arg \max_{\alpha^k} w^k(\alpha^k, r_k^{-k}(\rho^{-k})) - C^k(\alpha^k, r_k^{-k}(\rho^{-k}))]$ is UHC. Existence of perturbed equilibria then follows from the Kakutani fixed point theorem. \square

Theorem 12. *If ρ is justifiable then it is a collusion constrained equilibrium.*

Proof. Suppose ρ is justifiable. Then there exists a sequence of perturbations $r_{kn}^{-k}, C_n^k, w_n^k$ such that $\max_{\rho^{-k}} |r_{kn}^{-k}(\rho^{-k}) - \rho^{-k}| \rightarrow 0$, C_n^k is a high cost sequence, and $\max_{\alpha^k, \rho^{-k}} |w_n^k(\alpha^k, \rho^{-k}) - v^k(\alpha^k, \rho^{-k})| \rightarrow 0$, each with a perturbed equilibrium ρ_n that converges to ρ .

Let $B_{ucn}^k = \arg \max_{\alpha^k} w_n^k(\alpha^k, r_{kn}^{-k}(\rho^{-k})) - C_n^k(\alpha^k, r_{kn}^{-k}(\rho^{-k}))$. Let $\tilde{v} = \max v^k - \min v^k$. Let $\delta_{n1} = \max_{\alpha^k, \rho^{-k}} |w_n^k(\alpha^k, r_{kn}^{-k}(\rho^{-k})) - w_n^k(\alpha^k, \rho^{-k})|$ and $\delta_{n2} = \max_{\alpha^k, \rho^{-k}} |w_n^k(\alpha^k, \rho^{-k}) - v^k(\alpha^k, \rho^{-k})|$. Since C_n^k is a high cost sequence, for all large enough n , $G^k(\alpha^k, \rho^{-k}) > \gamma_n^k$

would imply $C_n^k(\alpha^k, \rho^{-k}) > 2(\tilde{v} + \delta_{n1} + \delta_{n2})$ and since $\max_{\rho^{-k}} |r_{kn}^{-k}(\rho^{-k}) - \rho^{-k}| \rightarrow 0$, also $C_n^k(\alpha^k, r_{kn}^{-k}(\rho^{-k})) > \tilde{v} + \delta_{n1} + \delta_{n2}$. So for all sufficiently large n , $\alpha^k \in B_{wcn}^k$ would mean $G^k(\alpha^k, \rho^{-k}) \leq \gamma_n^k$.

Let $W_n^k(\rho^{-k}) = \max_{\alpha^k \in \mathcal{A}^k | G^k(\alpha^k, r_{kn}^{-k}(\rho^{-k}))=0} w_n^k(\alpha^k, r_{kn}^{-k}(\rho^{-k}))$. Suppose $\alpha_n^k \in B_{wcn}^k$; then for large enough n it must be that

$$w_n^k(\alpha_n^k, r_{kn}^{-k}(\rho^{-k})) \geq W_n^k(\rho^{-k}) \geq V_S^k(\rho^{-k}) - \delta_{n1} - \delta_{n2}$$

This in turn means

$$v^k(\alpha_n^k, \rho^{-k}) \geq W_n^k(\rho^{-k}) - \delta_{n1} - \delta_{n2} \geq V_S^k(\rho^{-k}) - 2\delta_{n1} - 2\delta_{n2}$$

Notice that the sets B_{wcn}^k therefore satisfy the premise of Lemma 1 if we set the scalars V_n^k equal to $W_n^k(\rho^{-k}) - \delta_{n1} - \delta_{n2}$. So we know that ρ must be such that for all k , $\rho^k \in H(B^k)$ where B^k is the set of α^k that satisfies $G^k(\alpha^k, \rho^{-k}) = 0$ and $v^k(\alpha^k, \rho^{-k}) \geq \liminf V_n^k$. Finally note that

$$\liminf W_n^k(\rho^{-k}) - \delta_{n1} - \delta_{n2} \geq \liminf V_S^k(\rho^{-k}) - 2\delta_{n1} - 2\delta_{n2} \Rightarrow \liminf W_n^k(\rho^{-k}) \geq V_S^k(\rho^{-k}).$$

ρ is therefore a collusion constrained equilibrium. \square

Perturbed Equilibrium: Lower HemiContinuity

Theorem 13. *If ρ is a collusion constrained equilibrium then it is justifiable.*

Proof. We are given a collusion constrained equilibrium ρ and want to find a sequence of perturbations with perturbed equilibria $\rho_n \rightarrow \rho$. In fact the construction we are going to suggest will do something stronger, the idea is to construct a series of perturbations with perturbed equilibria $\rho_n = \rho$ which obviously converges to itself. Recall that $\rho^k \in H[B_S^k(\rho^{-k})]$. The idea is to find a perturbed equilibrium so that $\arg \max_{\alpha^k} w_n^k(\alpha^k, r_k^{-k}(\rho^{-k})) - C_n^k(\alpha^k, r_k^{-k}(\rho^{-k})) = B_S^k(\rho^{-k})$; then clearly ρ^k itself is in $H[\arg \max_{\alpha^k} w_n^k(\alpha^k, r_k^{-k}(\rho^{-k})) - C_n^k(\alpha^k, r_k^{-k}(\rho^{-k}))]$.

Step 1: Choose, for each k , a sequence σ_{kn}^{-k} with $\sigma_{kn}^{-k} \rightarrow \rho^{-k}$ and $V^k(\sigma_{kn}^{-k}) \rightarrow V_S^k(\rho^{-k})$. We know that we can find such a sequence by the definition of $V_S^k(\rho^{-k})$: it is the limit of the worst of the local best, so there must be some sequence of local bests that converges to it.

Constants: Define $\bar{G}^k(\sigma^{-k}) = \max_{\alpha^k} |G^k(\alpha^k, \sigma^{-k}) - G^k(\alpha^k, \rho^{-k})|$, $\bar{G}_n^k = \bar{G}^k(\sigma_{kn}^{-k})$, and similarly $\bar{V}(\sigma^{-k}) = \max\{0, V^k(\sigma^{-k}) - V_S^k(\rho^{-k})\}$, $\bar{V}_n^k = \bar{V}(\sigma_{kn}^{-k})$ and note that both \bar{G}_n^k and \bar{V}_n^k go to zero as $n \rightarrow \infty$. Also let $\bar{v}^k(\sigma^{-k}) = \max_{\alpha^k} |v^k(\alpha^k, \sigma^{-k}) - v^k(\alpha^k, \rho^{-k})|$ and $\bar{v}_n^k = \bar{v}^k(\sigma_{kn}^{-k})$; observe that $\bar{v}_n^k \rightarrow 0$. Take $\lambda_n^k = 1/\sqrt{\bar{G}_n^k}$ which goes to infinity, $\bar{\kappa}_n^k = 3(\bar{v}_n^k + \bar{V}_n^k + \lambda_n^k \bar{G}_n^k)$ which goes to zero and $\bar{\gamma}_n^k = 1/\sqrt{\lambda_n^k}$ which goes to zero.

The functions $\bar{w}_n^k(\alpha^k, \sigma^{-k})$ and $\bar{C}_n^k(\alpha^k, \sigma^{-k})$: Define first $D_n^k(\alpha^k) = \max\{0, v^k(\alpha^k, \rho^{-k}) - V_S^k(\rho^{-k})\} + \lambda_n^k G(\alpha^k, \rho^{-k})$ and $d_n^k(\alpha^k) = \min\{D_n^k(\alpha^k), \bar{\kappa}_n^k\}$. This converges uniformly to zero. We then take $\bar{C}_n^k(\alpha^k, \sigma^{-k}) = D_n^k(\alpha^k) - d_n^k(\alpha^k)$ and $\bar{w}_n^k(\alpha^k, \sigma^{-k}) = v^k(\alpha^k, \rho^{-k}) - d_n^k(\alpha^k)$. Observe that

$$\begin{aligned} \bar{w}_n^k(\alpha^k, \sigma^{-k}) - \bar{C}_n^k(\alpha^k, \sigma^{-k}) &= v^k(\alpha^k, \rho^{-k}) - D_n^k(\alpha^k) \\ &= v^k(\alpha^k, \rho^{-k}) - \max\{0, v^k(\alpha^k, \rho^{-k}) - V_S^k(\rho^{-k})\} - \lambda_n^k G(\alpha^k, \rho^{-k}) \\ &= \min\{v^k(\alpha^k, \rho^{-k}), V_S^k(\rho^{-k})\} - \lambda_n^k G(\alpha^k, \rho^{-k}) \end{aligned}$$

Key fact: $\arg \max_{\alpha^k} \bar{w}_n^k(\alpha^k, \sigma^{-k}) - \bar{C}_n^k(\alpha^k, \sigma^{-k}) = B_S^k(\rho^{-k})$. To see this consider the maximizers of $\min\{v^k(\alpha^k, \rho^{-k}), V_S^k(\rho^{-k})\} - \lambda_n^k G(\alpha^k, \rho^{-k})$. For the elements of $B_S^k(\rho^{-k})$, that is the α^k for which $G(\alpha^k, \rho^{-k}) = 0$ and $v^k(\alpha^k, \rho^{-k}) \geq V_S^k(\rho^{-k})$, the expression equals $V_S^k(\rho^{-k})$. Outside $B_S^k(\rho^{-k})$, that is for α^k such that $G^k(\alpha^k, \rho^{-k}) > 0$ or $v^k(\alpha^k, \rho^{-k}) < V_S^k(\rho^{-k})$, the expression is lower than that value. This proves the assertion.

Properties: There exists $\epsilon_n^k > 0$ such that $|\sigma^{-k} - \sigma_{kn}^{-k}| \leq \epsilon_n^k$ implies

- (i) if $G^k(\alpha^k, \sigma^{-k}) > \bar{\gamma}_n^k$ then $\bar{C}_n^k(\alpha^k, \sigma^{-k}) \geq \lambda_n^k \bar{\gamma}_n^k - \bar{\kappa}_n^k - 2\lambda_n^k \bar{G}_n^k \rightarrow \infty$
- (ii) if $G^k(\alpha^k, \sigma^{-k}) = 0$ then $\bar{C}_n^k(\alpha^k, \sigma^{-k}) = 0$
- (iii) $|\bar{w}_n^k(\alpha^k, \sigma^{-k}) - v^k(\alpha^k, \sigma^{-k})| \leq 2\bar{v}_n^k + \bar{\kappa}_n^k \rightarrow 0$

Proof of these:

(i) $\bar{C}_n^k(\alpha^k, \sigma^{-k}) \geq \lambda_n^k G(\alpha^k, \rho^{-k}) - \bar{\kappa}_n^k \geq \lambda_n^k G(\alpha^k, \sigma^{-k}) - \bar{\kappa}_n^k - \lambda_n^k \bar{G}_n^k(\sigma^{-k})$, so choose ϵ_n^k small enough that $\bar{G}_n^k(\sigma^{-k}) \leq 2\bar{G}_n^k$.

(ii) Choose $\epsilon_n^k > 0$ such that for all $|\sigma^{-k} - \sigma_{kn}^{-k}| < \epsilon_n^k$ we have $\max_{\alpha^k} |G^k(\alpha^k, \sigma^{-k}) - G^k(\alpha^k, \sigma_{kn}^{-k})| \leq \bar{G}_n^k$. Note that $\max_{\alpha^k} |G^k(\alpha^k, \sigma_{kn}^{-k}) - G^k(\alpha^k, \rho^{-k})| = \bar{G}_n^k$. Hence by the triangle inequality $G^k(\alpha^k, \sigma^{-k}) = 0$ implies $G^k(\alpha^k, \rho^{-k}) \leq 2\bar{G}_n^k$.

Since V_S^k can not jump up we may choose $\epsilon_n^k > 0$ such that for all $|\sigma^{-k} - \sigma_{kn}^{-k}| < \epsilon_n^k$ we have $V_S^k(\sigma^{-k}) \leq V_S^k(\sigma_{kn}^{-k}) + \bar{v}_n^k$. Note that $V_S^k(\sigma_{kn}^{-k}) \leq V_S^k(\rho^{-k}) + \bar{V}_n^k$. Hence $V_S^k(\sigma^{-k}) \leq V_S^k(\rho^{-k}) + \bar{v}_n^k + \bar{V}_n^k$. Therefore $G^k(\alpha^k, \sigma^{-k}) = 0$ implies $v^k(\alpha^k, \sigma^{-k}) \leq V_S^k(\rho^{-k}) + \bar{v}_n^k + \bar{V}_n^k$.

Finally choose $\epsilon_n^k > 0$ such that for all $|\sigma^{-k} - \sigma_{kn}^{-k}| < \epsilon_n^k$ we have $\max_{\alpha^k} |v^k(\alpha^k, \sigma^{-k}) - v^k(\alpha^k, \sigma_{kn}^{-k})| \leq \bar{v}_n^k$. Hence by the triangle inequality $\max_{\alpha^k} |v^k(\alpha^k, \sigma^{-k}) - v^k(\alpha^k, \rho^{-k})| \leq 2\bar{v}_n^k$.

Putting these inequalities together we see that $G^k(\alpha^k, \sigma^{-k}) = 0$ implies that $D_n^k(\alpha^k) = \max\{0, v^k(\alpha^k, \rho^{-k}) - V_S^k(\rho^{-k})\} + \lambda_n^k G(\alpha^k, \rho^{-k}) \leq 3\bar{v}_n^k + \bar{V}_n^k + 2\lambda_n^k \bar{G}_n^k \leq \bar{\kappa}_n^k$, which in turn implies $\bar{C}_n^k(\alpha^k, \sigma^{-k}) = 0$.

(iii) Recalling that $\epsilon_n^k > 0$ is such that for all $|\sigma^{-k} - \sigma_{kn}^{-k}| < \epsilon_n^k$ we have $\max_{\alpha^k} |v^k(\alpha^k, \sigma^{-k}) - v^k(\alpha^k, \sigma_{kn}^{-k})| \leq \bar{v}_n^k$, property (iii) follows from

$$\begin{aligned} &|\bar{w}_n^k(\alpha^k, \sigma^{-k}) - v^k(\alpha^k, \sigma^{-k})| \\ &\leq |v^k(\alpha^k, \sigma^{-k}) - v^k(\alpha^k, \sigma_{kn}^{-k})| + |v^k(\alpha^k, \sigma_{kn}^{-k}) - v^k(\alpha^k, \rho^{-k})| + d_n^k(\alpha^k) \leq 2\bar{v}_n^k + \bar{\kappa}_n^k \end{aligned}$$

Step 2: We now have $\bar{w}_n^k(\alpha^k, \sigma^{-k})$ and $\bar{C}_n^k(\alpha^k, \sigma^{-k})$ which are defined in a ϵ_n^k -neighborhood of σ_{kn}^{-k} and have the right properties there. For $|\sigma^{-k} - \rho^{-k}| < \epsilon_n^k$ we define $\bar{r}_{kn}^{-k}(\sigma^{-k}) = \sigma_{kn}^{-k}$ (taking advantage of the fact that these need not be the same for all k). We must now extend these to functions $w_n^k(\alpha^k, \sigma^{-k})$, $C_n^k(\alpha^k, \sigma^{-k})$, $r_{kn}^{-k}(\sigma^{-k})$ on all of R^{-k} while preserving the right properties and the values of $\bar{w}_n^k(\alpha^k, \sigma_{kn}^{-k})$, $\bar{C}_n^k(\alpha^k, \sigma_{kn}^{-k})$ and $\bar{r}_{kn}^{-k}(\rho^{-k})$. We can do this with a simple pasting. Let $\beta_n^k(x)$ be a non-negative continuous real valued function taking the value of 1 at $x = 0$ and the value of 0 for $x \geq \epsilon_n^k$. Then we define

$$\begin{aligned} w_n^k(\alpha^k, \sigma^{-k}) &= \beta_n^k(|\sigma^{-k} - \sigma_{kn}^{-k}|) \bar{w}_n^k(\alpha^k, \sigma^{-k}) + (1 - \beta_n^k(|\sigma^{-k} - \sigma_{kn}^{-k}|)) v^k(\alpha^k, \sigma^{-k}) \\ C_n^k(\alpha^k, \sigma^{-k}) &= \beta_n^k(|\sigma^{-k} - \sigma_{kn}^{-k}|) \bar{C}_n^k(\alpha^k, \sigma^{-k}) + (1 - \beta_n^k(|\sigma^{-k} - \sigma_{kn}^{-k}|)) \lambda_n^k G^k(\alpha^k, \sigma^{-k}) \\ r_n^k(\sigma^{-k}) &= \beta_n^k(|\sigma^{-k} - \sigma_{kn}^{-k}|) \bar{r}_n^k(\sigma^{-k}) + (1 - \beta_n^k(|\sigma^{-k} - \sigma_{kn}^{-k}|)) \sigma^{-k}. \end{aligned}$$

It is easy to check that these pasted functions have the correct properties. Note that requiring $\bar{w}_n^k(\alpha^k, \sigma^{-k})$ and $\bar{C}_n^k(\alpha^k, \sigma^{-k})$ to have the right properties in the ϵ_n^k -neighborhood of σ_{kn}^{-k} ensures that the above convex combinations inherit those properties. \square