Bilateral Treaties*

Sophie Bade[†] Guillaume Haeringer[‡]
Ludovic Renou[§]

4th November 2005

*We are grateful to Doh-Shin Jeon, Efe Ok, seminar participants at Adelaide, Barcelona, UNSW, Sydney, Paris and the Bocconi University in Milan for helpful comments and suggestions. We would like to warmly thank John van der Hoek and Guillaume Carlier for patiently and brilliantly answering our questions on differential topology and the theory of non-smooth optimizations. Haeringer acknowledges a financial support from the Spanish Ministry of Science and Technology through grant BEC2002- 02130 and from the Barcelona Economics - CREA program. Renou acknowledges financial support from the Faculty of the Professions.

[†]Department of Economics, Penn State University, 518 Kern Graduate Building, University Park, PA 16802-3306, USA. Phone: (814) 865 8871, Fax: (814) 863 - 4775 `sub18@psu.edu`

[‡]Departament d'Economia i d'Història Econòmica, Universitat Autònoma de Barcelona, 08193 Bellaterra, Spain. Phone: (+34) 93 581 12 15, Fax: (+34) 93 581 20 12. `Guillaume.Haeringer@uab.es`

[§]Department of Economics, Room G52, Napier Building, Adelaide 5005, Australia. Phone: +61 (0)8 8303 4930, Fax: +61 (0)8 8223 1460. `ludovic.renou@adelaide.edu.au`

Abstract

We consider purely non-cooperative environments, in which two parties interact. We assume that the parties cannot sign binding agreements, but have the power to restrict their own action sets. Such commitments by each party are called *treaties*. We completely characterize the set of action profiles that can be implemented by a treaty. We then show that implementable profiles are generically inefficient. We are nonetheless able to show that treaties can improve upon the status quo, that is, the situation without commitment.

KEYWORDS: Commitment, self-enforcing, treaties, inefficiency, agreements, Pareto-improvement.

JEL CLASSIFICATION NUMBERS: C7.

1 Introduction

Pollution control, disarmament, capacity choices for firms, or importation quotas are, just to name a few, situations in which agents may have an interest to refrain from choosing some actions. When commitment to reduce one's action set is irreversible such a decision may not be profitable if taken alone, however. A typical example is disarmament between two countries, where an overall reduction of ammunitions, though desirable, is difficult to obtain if it is done by one country only. No country is willing to do it if the other does not.[1] Coordination on which commitment to make could then be achieved by signing a "contract" that specifies which action each party is supposed to take. But this type of agreement becomes difficult to sustain when no third party can impose penalties on defectors, or more generally if the environment is purely non-cooperative. In this case a *joint* agreement has no intrinsic value and in the best case it could be considered as a mere cheap talk.[2] A possible solution is that parties police each other, meaning that an agreement should also specify what actions each party is allowed to undertake when an opponent defects. In this respect an agreement may not necessarily specify a single action, but a *set* of actions. We call such agreements *treaties*. Committing to a set of actions and not a single action makes the opponents loosing control on which action will be chosen, however, and therefore legitimates deviations from the outcome to be implemented. The design of a treaty is then not a trivial task: One the one hand we should allow parties to commit to a set of actions (rather than a single action) so as to give them means to threat possible deviators but on the other hand, theses sets of actions should not tempt parties to deviate from the action profile the treaty is aimed to implement.

Game theoretically speaking, a treaty specifies a restriction of the strategy set available to each player (including, possibly, not restricting at all). As a preliminary observation, we want to stress that a restriction of strategy sets

[1]Gradualism in the terms of commitment is often invoked to overcome incentives to free-ride. See for instance Lockwood and Thomas (2002) for a repeated Prisoners' Dilemma situation where players can gradually increase their level of cooperation.

[2]See, for instance, Ray and Vohra (1997) for situations where groups of agents can make binding decisions.

amounts to define a new game, the *induced game*, where payoffs are entirely obtained from the original game (any strategy profile of the new game is also a strategy profile of the original game). A treaty is self-enforcing when no player has an incentive to take a different commitment than the one under consideration, i.e., no player would take a different restriction of her strategy set. Since we shall assume that once players have made their commitments they play a Nash equilibrium of the induced game, a treaty is then self-enforcing if it is the first stage outcome of a a subgame perfect equilibrium of the following two-stage game: First, players choose a restriction of the action sets, and second they play the induced game. Therefore, self-enforceability of a treaty is *mutual* and is obtained through a simple sequential game structure, without assuming punishment scheme against deviating players. We then call a profile *implementable by a treaty* if it is a Nash equilibrium of the game induced by a self-enforcing treaty. [3]

Our goal in this paper is to characterize the set of profiles implementable by a treaty. There are two main issues in this characterization. First we derive a simple procedure to check whether a profile is implementable. Second, we analyze the welfare properties of implementable profiles.

Since we assume in this paper that players' action sets are subset of the real line, for action profile x there is an infinite number of restrictions of the original strategy set that contain the profile x. Any of these sets can lead to this profile, which makes the characterization of implementable profiles not a trivial task. In spite of this apparent difficulty we are able to offer a complete characterization of implementable profiles when there are two players. This characterization turns out to produce a very simple method

[3]Maggi and Morelli (2004) and Conconi and Perroni (2005) also point out that international agreements must be self-enforcing because countries are sovereign and therefore collective agreements cannot be binding for participants. Conconi and Perroni liken the ability of a country to make a binding decision to domestic credibility. Their approach differs from ours, however, as they consider the trade-off between international cooperation and internal policy making. Their main finding is that deviations from an international agreement can become less attractive when countries are unable to commit domestically. Conconi and Perroni thus argue that the ability to commit may hinder international cooperation. Maggi and Morelli consider voting issues in international organization and show that unanimity is the optimal rule if one wants voting decisions to be self-enforcing.

to check whether a profile is implementable. Namely, we show that a profile is implementable by a treaty if and only if it is implementable by a certain standardized form of treaties, which we call *simple treaties*. Such treaties are defined by one player restricting his strategy set to a singleton and for the other player taking a restriction that contains his best-reply to this action.[4] The advantage of such treaties is twofold. First, any action profile can be obtained by at most four simple treaties, thus reducing significantly the set of treaties candidate to implement a profile. Second, simple treaties reduce the set of potential deviators to one player only, and with further scrutiny we can reduce the set of candidate treaties to be only two. This is so because we show that a profile is implementable by a treaty if and only if at least one player is playing a best-reply to the opponent's action.

This paper is organized as follows. In the next section, we develop an example that previews our approach and our main results. In Section 3, we give a detailed description of the environment faced by the parties, and define what we call the game of treaties. Section 4 presents some preliminary results. We then completely characterize in Section 5 the set of action profiles that are implementable by self-enforcing treaties. Section 6 discusses the welfare implications of self-enforcing treaties. Proofs are relegated in the Appendix.

2 An example

Two countries, 1 and 2, compete on the international market for the production of *widgets*. To produce *widgets*, two chemical products C and D can be used. The product C is cheap but pollutes heavily, while product D is more environmentally friendly. We denote by x_i the fraction of product D used in country i, $i \in \{1, 2\}$.[5] The cost of using product C is normalized to zero. The cost of using product D is equal to $\frac{1}{2}x_i^2$ for a level of x_i used in country i. Moreover, there are network effects: the more a country is using product D, the higher the benefit for the other country to use product D as well. For country i, the benefit for using a level x_i of product D is $(x_i/2)\,(5/2 - x_i/x_j)$

[4] The exact definition adds other requirements with respect to the bounds of the restriction of this second player.

[5] For technical reasons, we assume that $x_i \in [\varepsilon, 1]$ for $\varepsilon > 0$, but small enough.

if country $j \neq i$ uses a level of x_j. The payoff of country i if it uses $x_i\%$ of product D and country j uses $x_j\%$ is therefore

$$u_i(x_i, x_j) = \frac{1}{2}x_i\left(\frac{5}{2} - \frac{x_i}{x_j}\right) - \frac{1}{2}x_i^2,\ i \neq j.$$

If the government of each country does not impose restrictions on the proportion of product D being used, each country uses 25% of product D in equilibrium i.e., $(\frac{1}{4}, \frac{1}{4})$ is the unique equilibrium of the game described above. We will refer to this game as *the mother game*. For future reference, note that the (mother) best-reply of country i to country j using $x_j\%$ of product D is

$$BR_i(x_j) = \frac{5x_j}{4(1+x_j)}.$$

Suppose that before deciding how much to use of product D, the government of each country can impose some restrictions on the use of the heavy pollutant C.[6,7] In this paper, we call a pair of restrictions (X_1, X_2) with $X_i = [\underline{x}_i, \overline{x}_i]$, $i \in \{1, 2\}$, a *bilateral treaty*. Choosing first restrictions of action sets and then choosing an action (from within these restrictions) defines a two-stage game that we call *game of treaties*.

Observe that any bilateral treaty (X_1, X_2) induces a game with the payoff function u_i, but a restricted set of actions $X_i \subseteq [0, 1]$ for each country. The central theme of this paper is the complete characterization of the set of action profiles x that can be implemented by a treaty. We say that an action profile is *implementable by a treaty* if it is played on the equilibrium path in the game of treaties.

---

[6] We assume that the preferences of the government agree with the preferences of the country.

[7] For instance, a government can impose the exclusive use of the more environmental friendly product D, or that at least $\underline{x}\%$ of product D has to be used. A government can thus impose a lower bound on the use of product D. Similarly, a government might want to impose an upper bound $\overline{x}$ on the use of product D e.g., a country's competition policy might require both products C and D to be used or a country might apply a precautionary principle avoiding the use of a single product.

A first result is that a treaty can always perpetuate a preexisting state of affairs, that is, Nash equilibria of the game without restrictions (the *mother game*) are implementable. So, how to implement $(1/4, 1/4)$? It is easy to see that the equilibrium $(1/4, 1/4)$ is implemented by the bilateral treaty $(\{1/4\}, \{1/4\})$. If each country commits to 1/4 in the first-stage of the game of treaties, the induced game has the unique action 1/4 per player, hence $(1/4, 1/4)$ is an equilibrium of this induced game. Moreover, it is easy to see that given the commitment of one player to $\{1/4\}$ in the first-stage, the other player cannot profitably deviate in the first-stage as 1/4 is his best-reply to 1/4. (Off equilibrium path, both countries play any equilibrium of the induced game.)

Besides implementing the equilibrium $(1/4, 1/4)$, treaties can implement a wide range of other action profiles. For instance, another profile that is implementable is $x^* = (2/5, 8/17)$. A pair of commitments supporting the profile x^* is $X_1 = [0, 2/5]$ and $X_2 = \{8/17\}$, respectively. To see this, note that since country 1 commits to X_1, country 1's "new" best-reply br_1 is,[8]

$$br_1(x_2) = \begin{cases} 2/5 & \text{if } BR_1(x_2) \leq 2/5, \\ BR_1(q_2) & \text{if } BR_1(x_2) > 2/5. \end{cases}$$

Note also that $(2/5, 8/17)$ is an equilibrium of the game induced by the commitment of each country to $X_1 = [0, 2/5]$ and $X_2 = \{8/17\}$, respectively. Moreover, since 2/5 is the (mother) best-reply $BR_1(8/17)$ of player 1 to the commitment 8/17 of country 2, country 1 cannot profitably deviate from his commitment X_1.[9] Let us now consider country 2. Since in any game induced by a deviation of country 2, countries play an equilibrium of this induced game, it follows that country 2 has a profitable deviation only if there exists a level x'_2 such that $u_2(x'_2, br_1(x'_2)) > u_2(8/17, 2/5)$. As $br_1(8/17) = 2/5$, this is equivalent to $u_2(x'_2, br_1(x'_2)) > u_2(8/17, br_1(8/17))$. However, the unique maximizer of $u_2(\cdot, br_1(\cdot))$ is 8/17, hence country 2 has no profitable deviation in the first stage.

Observe that the treaty we use to implement the profile $(2/5, 8/17)$ has two distinctive features. First, country 2 commits to a single action 8/17.

[8]See Lemma 1 in Section 4.

[9]Off equilibrium path, let both countries play any equilibrium of the induced game.

Second, country 1 only truncates its set of actions at the top, and, moreover, the truncation is made at $BR_1(8/17) = 2/5$, the (mother) best-reply of country 1 to the extreme commitment of country 2. We call any treaty with these distinctive features a *simple treaty*. It turns out that any implementable profile of actions is implementable by a simple treaty (see Theorem 1). Furthermore, it is easy to see that this implies that the set of implementable profiles of actions is necessarily a subset of the graphs of both players' best-replies. In this example, the set of implementable profiles of actions is

$$\begin{aligned} &\{x : x_1 = BR_1(x_2),\ 1/4 \leq x_2 \leq 17/36\} \bigcup \\ &\{x : x_2 = BR_1(x_1),\ 1/4 \leq x_1 \leq 17/36\}. \end{aligned} \tag{1}$$

For instance, it is easy to check that the profile $(17/36, 85/202)$ is implementable by the simple treaty $(\{17/36\}, [0, 85/202])$.

Two additional remarks are in order. First, in any implementable profile, both countries use more of the environmental friendly product D than in the Nash equilibrium $(1/4, 1/4)$. Paradoxically, countries need to impose a maximum percentage on the use of product D to implement such profiles e.g., country 1 need to impose an upper bound of $2/5$ on the use of D to implement the profile $(2/5, 8/17)$. This result is rather counter-intuitive as we might expect that in order to use more of product D, countries need to impose minimum requirements rather than maximum requirements. Second, it is worth noting that the profile $(2/5, 8/17)$ makes both countries better off compared to the Nash payoffs. Therefore, a treaty can implement a Pareto-improvement upon the mother game.

Finally, we can observe that none of the implementable profiles described in (1) is *efficient*. For instance, the profile that maximizes the sum of the welfare of the countries, $(3/4, 3/4)$ does not belong to the set of implementable profiles. This observation will be confirmed in our second main result (Theorem 2) which states that implementable profiles are generically inefficient.

3 Games of treaties

3.1 Preliminaries

The initial situation we consider is a two-player game $G := \langle N, (Y_i, u_i)_{i\in N}\rangle$ with $N = \{1, 2\}$ the set of players, Y_i the set of actions available to player i, and $u_i : Y_1 \times Y_2 \to \mathbb{R}$ the payoff function of player i. Denote $Y := Y_1 \times Y_2$. We call the opponent of player i, player j. We assume that for each player $i \in \{1, 2\}$, Y_i is a non-empty, compact, convex subset of the real line. Without loss of generality, we take $Y_i = [0, 1]$, for $i \in \{1, 2\}$. For each player i, the payoff function u_i is assumed to be continuous in all its arguments and strictly quasi-concave in y_i, i.e., for all $y_j \in Y_j$, $y_i \in Y_i$, $y_i' \in Y_i$ and $\alpha \in (0, 1)$, $u_i(\alpha y_i + (1 - \alpha)y_i', y_j) > \min(u_i(y_i, y_j), u_i(y_i', y_j))$.[10]

We furthermore assume that players have the ability to unilaterally commit not to play some actions, i.e., to restrict their action sets. Such commitments are assumed to be perfectly binding, meaning that if player i restricts his action set to X_i, any action chosen later on must belong to X_i.[11]

Definition 1 A *(bilateral) treaty* is a pair (X_1, X_2) where X_i is a non-empty, compact and convex subset Y_i, for all $i \in \{1, 2\}$.

Thus, our definition of a treaty imposes on each player a *restriction* of their action sets. A treaty does not prescribe the choice of an action profile once each player has made their choice about which actions to discard. In the words of Hart and Moore (2004), "in a bilateral treaty, the parties commit not to consider actions not on the list (X_1, X_2), i.e., these actions are ruled out. Ex-post, the parties are free to choose from the list of actions specified in the treaty i.e., actions are not ruled in."

Remark 1 Equivalently, we can write the restricted action space X_i of player i as a closed real interval $[\underline{x}_i, \overline{x}_i] \subseteq [0, 1]$, where $\underline{x}_i$ ($\overline{x}_i$) refers to the

[10]In the words of Moulin (1984), G is a two-player *nice* game.

[11]Why or how parties can commit is an issue we will not address. Our purpose here is to see how the ability to make binding commitment enables players to make self-enforcing agreements. Admati and Perry (1991) study the problem of voluntary contributions to a joint project and look at different regimes of commitment. See Bénabou and Tirole (2004) for the problem of personal commitment.

minimum (maximum) of player i's restricted action space. Note that player i can also commit to a singleton, in which case $\underline{x}_i = \overline{x}_i$.

We say that the bilateral treaty (X_1, X_2) *induces* the game $G(X) := \langle N, (X_i, u_i^X) \rangle$, where $u_i^X(x) = u_i(x)$ for all $x \in X$, $i \in \{1,2\}$. Abusing notation, we will drop the superscript X in the sequel. The *induced* game $G(X)$ is thus obtained from the game G by restricting the action sets of the players. We shall use the term "*mother*" to make reference to to the original game G. For instance, we shall use the expressions *mother game*, *mother best-reply*, *mother action set*, etc. Similarly, the term "*induced*" will refer the best reply, action sets etc. in $G(X)$. We denote by $\mathcal{Y}_i$ the collection of all non-empty, compact, convex subsets of Y_i, and define $\mathcal{Y} := \prod_{i \in N} \mathcal{Y}_i$.

3.2 Games of treaties

Given the strategic-form game G, the *game of treaties* $\Gamma(G)$ is a two-stage game with perfect information, in which:

Stage 1. Both players simultaneously choose action sets $X_i \in \mathcal{Y}_i$.

Stage 2. Players play the induced game $G(X)$.

A *strategy* for a player i in the game $\Gamma(G)$ (for short, Γ), is a pair $s_i = (X_i, \sigma_i)$ where $X_i \in \mathcal{Y}_i$, and σ_i is a mapping from $\mathcal{Y}$ to Y_i such that $\sigma_i(X) \in X_i$, for all $X \in \mathcal{Y}$. That is, a strategy for a player prescribes a choice of a restriction X_i (first-stage action) and, for each possible choice of a restriction for *both* players in the first-stage, an action $x_i \in X_i$ (second-stage action). The *outcome* of a strategy profile $s = (s_i)_{i \in \{1,2\}}$ is the pair (X, x) where $x_i = \sigma_i(X)$ for each player $i \in N$. The utilities over outcomes (x, X), or *payoffs* are assumed to only depend on the action profiles chosen in the second stage of the game and are given by the payoffs of the induced game. That is, we assume that player i derives utility $u_i(x)$ from outcome (x, X). If (x, X) is the outcome of strategy profile s we call x the *result* of s.

The central concept of this paper is the concept of implementation by a treaty, which we now define.

Definition 2 An action profile x is *implementable by a treaty* X if the pair (X, x) is the outcome of a subgame-perfect equilibrium of Γ.

Hence, a profile x is implementable by a treaty if it is a (stage 2) result of a subgame-perfect equilibrium of Γ. In this paper, we focus on subgame-perfect equilibria in pure strategies.

4 Games induced by treaties

A complete characterization of all implementable action profiles thus requires us to study the subgame perfect equilibria of Γ. To do so, we first derive some results concerning the proper subgames of Γ, namely the set of all induced games $G(X)$. The proofs of the results presented below, Lemmata 1 and 2 are in our companion paper, Bade, Haeringer, and Renou (2005).

So to proceed, we define $BR_i : Y_j \to Y_i$, the *(mother) best-reply* of player i in the game G, with for $y_j \in Y_j$,

$$BR_i(y_j) = \{y_i \in Y_i : u_i(y_i, y_j) \geq u_i(y_i', y_j) \text{ for all } y_i' \in Y_i\}.$$

When players commit to play in the set X, the best-reply map $br_i^X : X_{-i} \to X_i$ of player i is defined similarly, bearing in mind that now player i cannot choose an action outside X_i, that is, for all $x_j \in X_j$,

$$br_i^X(x_j) = \{x_i \in X_i : u_i(x_i, x_j) \geq u_i(x_i', x_j) \text{ for all } x_i' \in X_i\}.$$

We will denote the best-reply map $br_i^{X_i \times [0,1]}$ by $br_i^{X_i}$. Note that best-reply maps are non-empty, single valued and continuous. The strict quasi-concavity of payoff functions enables us to easily characterize br_i^X as a function of BR_i.

Lemma 1 *Player i's best-reply function in $G(X)$, $br_i^X : X_j \to X_i$, is*

$$br_i^X(x_j) = \begin{cases} \underline{x}_i & \text{if } BR_i(x_j) < \underline{x}_i\,, \\ BR_i(x_{-i}) & \text{if } \underline{x}_i \leq BR_i(x_j) \leq \overline{x}_i\,, \\ \overline{x}_i & \text{if } \overline{x}_i < BR_i(x_j)\,. \end{cases}$$

In words, the best-reply map br_i^X of the restricted game $G(X)$ agrees with the best-reply map BR_i of the mother game G on the set $\{x_j \in X_j : BR_i(x_j) \in X_i\}$, and is either $\underline{x}_i$ or $\overline{x}_i$, otherwise. Lemma 1 is illustrated in Figures (1a) and (1b). In the former it display a mother best-reply of player j and in the latter the restricted best-reply when he commits to $[\underline{x}_j, \overline{x}_j]$.

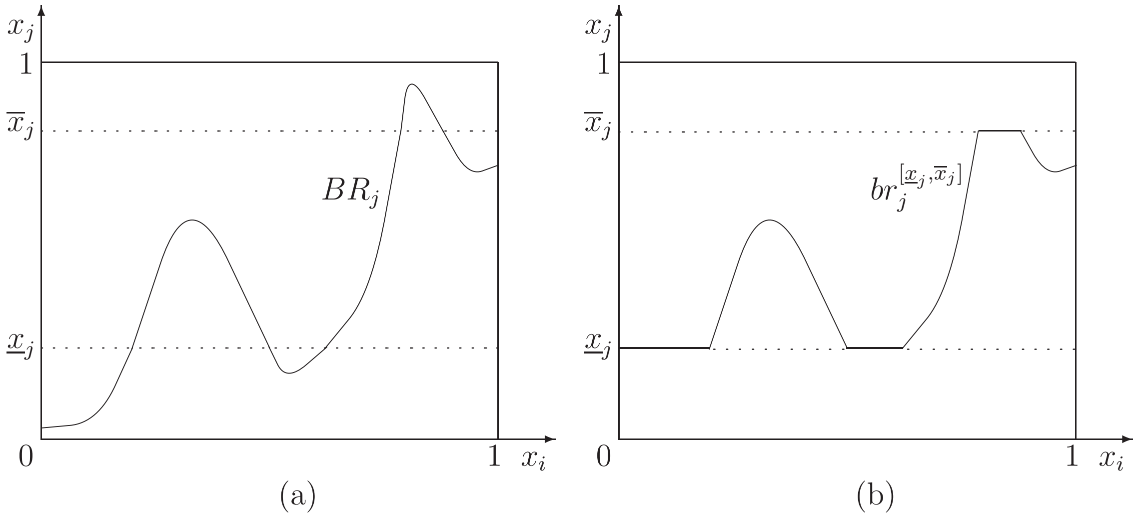

Figure 1: Mother and restricted best-replies

Denote $N(G)$ and $N(G(X))$ the set of Nash equilibria of G and $G(X)$, respectively. Observe that the mother game G as well as any induced game $G(X)$ has a Nash equilibrium in pure actions. Our next lemma states that if a profile of actions x^* is an equilibrium of $G(X)$, but is not an equilibrium of the mother game G, then $x^* \in \text{bd}_Y(X)$, the relative boundary of X in Y.

Lemma 2 *If $x^* \in N(G(X)) \setminus N(G)$, then $x^* \in \text{bd}_Y(X)$.*

Lemma 2 thus implies that if a treaty X^* implements a result x^* that is not an equilibrium of G, then it must be the case that for at least one player, say i, the action x_i^* is either the maximum or the minimum of X_i^*.

5 Implementation by treaties

5.1 Existence

We start by observing that the existence of a subgame-perfect equilibrium of Γ is not, *a priori*, guaranteed, for the cardinality of each player's strategy set in Γ is infinite. It turns out, however, that the issue of equilibrium existence in our case can be easily solved.[12]

Proposition 1 *The game of treaties has an equilibrium.*

Proof Since $\Gamma(G)$ is a finite horizon game, we can use the one-shot deviation property to check that a profile is an equilibrium —see Osborne and Rubinstein (1994, p. 103). Choose $y^* \in N(G)$ and consider for each player i the strategy $s_i^* = (\{y_i^*\}, \sigma_i^*)$, with $(\sigma_i^*(X))_{i\in\{1,2\}}$ a Nash equilibrium of $G(X)$ for any first-stage actions (treaty) X. By construction, no player can profitably change his second-stage action. It remains to be checked that no player has a profitable first-stage deviation. To this end, it suffices to observe that neither player can obtain a strictly higher payoff than $u_i(y^*)$, since $y_i^* = BR_i(y_j^*)$ for both $i \in \{1, 2\}$. So keeping the restriction $\{y_i^*\}$ of player i fixed player j cannot increase his utility by changing the restrictions on his action space. ■

The key observation in the proof of Proposition 1 is that any Nash equilibrium of the original situation G can be obtained as an equilibrium result of the game of treaties Γ. So, treaties have the power to perpetuate an existing situation. Moreover, it should be noted that uniqueness is clearly not guaranteed. For instance, if G has a multiplicity of equilibria, then we can already construct a multiplicity of subgame-perfect equilibria of Γ.

[12] See, for instance, Harris, Reny, and Robson (1995) for results on the existence of subgame-perfect equilibria for continuous games with almost perfect information. It is worth noting that the result holds independently of the number of players involved in the mother game G.

5.2 A Complete Characterization

We are now ready to characterize the set of all action profiles that can be implemented by a treaty. The main result of this section is that if a profile of actions x is implementable, then it is implementable by one of a very small number of treaties, those that we call *simple*.

DEFINITION 3 A treaty X is *simple* if it has the form $(\{x_i\}, [0, BR_j(x_i)])$ or $(\{x_i\}, [BR_j(x_i), 1])$.

In a simple treaty, one player takes an extreme position, that of excluding all but one action x_i. The other player, player j, truncates his action space at the best reply to the only action in player i's action space. He either truncates his action space from below or from above. We are now ready to formally state the main result of this section:

THEOREM 1 *An action profile x^* is implementable by a treaty if and only if it is implementable by a simple treaty.*

Before proving this characterization result, let us briefly comment on the implications of this theorem (see Section 5.5. for more on this). If we want to check whether a particular profile can be implemented by a treaty, we only need to check wether it can be implemented by a simple treaty. This is a very manageable task, as for any action profile there are exactly 4 simple treaties that could implement it. These treaties are:

$$([0, x_1], \{x_2\}), \quad ([x_1, 1], \{x_2\}),$$
$$(\{x_1\}, [0, x_2]), \quad (\{x_1\}, [x_2, 1])\,.$$

It is not difficult to check wether an action profile can be implemented by one of these four simple treaties. Indeed, to check whether x^* is implementable by $(\{x_1^*\}, [0, BR_2(x_1^*)])$, it suffices to check whether player 1 has an incentive to change his restricted action space. Observe that in the second stage, neither player has an incentive to deviate (player 2 is playing the mother best-reply to player 1's action, and player 1 does not have any choice). Furthermore, given that player 1 commits to $\{x_1^*\}$, player 2 does not have an incentive to alter the restrictions of his action space: the mother best-reply to x_1^* is already contained in $[0, BR_2(x_1^*)]$). Therefore, we only need to check whether player 1 has an incentive to deviate in the first stage of the game. Consequently, the action profile is an equilibrium if x_1^* solves the following optimization program:

$$\arg\max_{x_1 \in [0,1]} \quad u_1(x_1, br_2^{[0,BR_2(x_1^*)]}(x_1)) \,. \tag{2}$$

In Section 5.5, we take this optimization program as a starting point for a geometric characterization of implementable profiles.

5.3 Proof of Theorem 1

In this section, we present the main steps leading to Theorem 1 and give intuitions for these intermediate results. Detailed proofs can be found in the Appendix. We start by showing that if a result x^* is implementable, then for at least one player $i \in \{1, 2\}$, x_i^* is a best-reply to x_j^*.

Proposition 2 *Let x^* be implementable by some treaty X^*. Then $x_i^* = BR_i(x_j^*)$ for at least one player $i \in \{1, 2\}$.*

To see the intuition behind Proposition 2, suppose that a profile x^* is implementable by the treaty X^* such that neither player is using his mother best-reply. This means that for both players the restriction due to the treaty is binding, in the sense that both players would choose a different action if they had not restricted themselves. The continuity of the best replies implies that for all of player 2's actions in a sufficiently small interval $(x_2^* - \varepsilon, x_2^* + \varepsilon)$ around x_2^*, player 1's restricted best reply remains x_1^*. Let us now consider a different restriction for player 2. Take a $\{x_2'\}$ such that x_2' is a) closer to his own best reply to x_1^* and b) inside the interval for which player 1 does not change his action in the second stage. The strict quasi-concavity of player 2's payoff function implies that the result (x_1^*, x_2') which is closer to his mother best reply is strictly preferred to x^*. (See Figure 2.)

Proposition 3 *Let x^* be implementable by some treaty X^*. Then x^* is also implementable by the treaty X', such that $X_i' = \{x_i^*\}$ and $X_j' = X_j^*$ for $x_j^* = BR_j(x_i^*)$.*

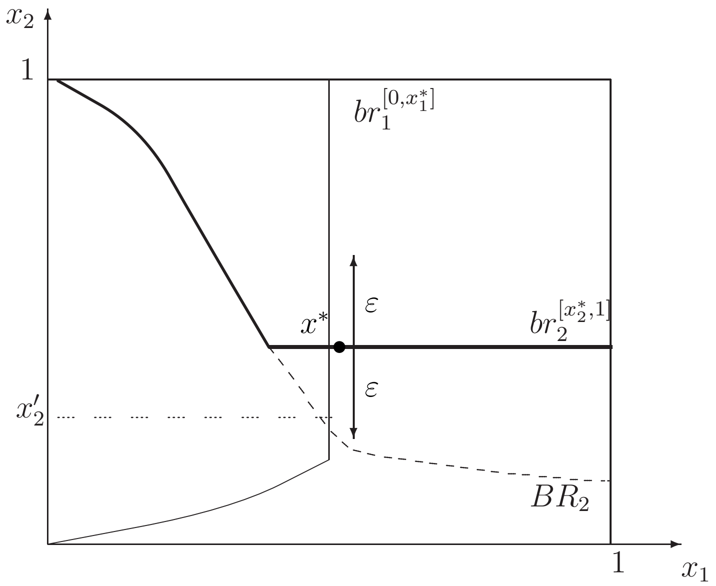

Figure 2: Illustration of Proposition 2

There is a tight connection between Proposition 2 and Proposition 3. By Proposition 2, we know that in any equilibrium outcome (X^*, x^*) of Γ, $x_j^* = BR_j(x_i^*)$ for player 1 or 2. Imagine now that player i commits to the singleton $\{x_i^*\}$. From the previous observation, player j has no incentive to deviate (he will still be able to play $BR_j(x_i^*)$ in the second stage). Hence, equilibrium payoffs and incentives to deviate are unchanged. If player i can profitably deviate when choosing the restriction $\{x_i^*\}$, he can also profitably deviate when choosing the restriction X_i^*. This, however, cannot be true as we started out with the assumption the (X^*, x^*) is an equilibrium outcome of the game.

The main insight of Proposition 3 is that if $(x_i^*, BR_j(x_i^*))$ is implementable by a treaty X^*, then the treaty X^* can be replaced by the treaty,

$$X' = (\{x_i^*\}, X_j^*). \tag{3}$$

To obtain Theorem 1, it suffices then to show that X_j^* can be reduced to be either $[0, x_j^*]$ or $[x_j^*, 1]$. We establish precisely that in the following proposition.

Proposition 4 *Let x^* be implementable by some treaty $(\{x_i^*\}, X_j^*)$. Then x^* is also implementable by a treaty X' such that $X'_i = \{x_i^*\}$ and either $X'_j = [BR_j(x_i^*), 1]$ or $X'_j = [0, BR_j(x_i^*)]$.*

Now to prove Theorem 1, take any implementable action profiles x^* and let X^* be a treaty that implements it. By Proposition 3, we know that a treaty $(\{x_i^*\}, X_j^*)$ for $i \in \{1, 2\}$ does also implement x^*. Finally, from Proposition 4, we know that an action profile that can be implemented by such a treaty can also be implemented by a simple treaty. In sum, these arguments imply that an action profile can be implemented by a treaty only if it can be implemented by a simple treaty. Conversely, any action profile that can be implemented by a simple treaty can be implemented by a treaty. This completes the proof of Theorem 1.

5.4 The geometry of implementable profiles

As already pointed out, Theorem 1 has remarkable implications for the characterization of the implementable action profiles of a game of treaties. To check whether a profile of actions x is implementable, it suffices to follow a simple four step procedure:

Step 1. Check whether x lies on the graph of the best-reply map of at least one player. If not, then x is not implementable. If yes, go to step 2.

Step 2. Check whether x lies on the best-reply graphs of both players. If yes, then x is implementable since it is an equilibrium of the mother game G. If not, go to step 3.

Step 3. Without loss of generality, assume that $x_j = BR_j(x_i)$. Construct the simple treaties $(\{x_i\}, [0, BR_j(x_i)])$ and $(\{x_i\}, [BR_j(x_i), 1])$. Go to step 4.

Step 4. Check whether x'_i maximizes $u_i(\cdot, br_j^{[0,BR_j(x_i)]}(\cdot))$ or $u_i(\cdot, br_j^{[BR_j(x_i),1]}(\cdot))$. If yes, then x is implementable. If not, then x is not implementable.

Steps 1 and 2 are easily translated into geometric analysis. An action profile can be implemented only if it lies on the best-reply curve of at least one player. If it lies on the best-reply curves of both players, this action profile is an equilibrium of the mother game, and from Proposition 1, it is implementable. Therefore, we are left with the question: *which of the action profiles that lie on only one best-reply curve can be implemented?* Steps 3 and 4 give the answer. However, these last two steps do not translate as easily into geometric analysis. In the sequel, we show that simple geometric arguments can be used to show that certain portions of the best-reply curves of the players *cannot* be implemented. Furthermore, we show that for a certain class of games, the set of implementable profiles can even be completely characterized by a straightforward geometric procedure.

To get this result, we first show that any equilibrium outcome can be described as a maximization program similar to that of Eq. (2).

PROPOSITION 5 *An outcome (X^*, x^*) is an equilibrium outcome if and only if, for at least one player $i \in \{1, 2\}$, $j \neq i$:*

(i) x_i^* *maximizes* $u_i(x_i, br_j^{X_j^*}(x_i))$ *, and*

(ii) $br_j^{X_j^*}(x_i^*) = BR_j(x_i^*)$.

Figure 3 illustrates the logic of Proposition 5. The outcome (x^*, X^*) with $X^* = (\{x_i^*\}, [0, \overline{x}_j])$ is an equilibrium outcome as the profile of actions x^* is associated with player i's highest indifference curves IC_i on the section of player j restricted best-reply curve $br_j^{[0,\overline{x}_j]}$ that corresponds with his mother best-reply curve BR_j. Observe that x^* is also implementable by the simple treaty $(\{x_i^*\}, [0, x_j^*])$, an illustration of Proposition 4.

REMARK 2 From Proposition 5, we have that x^* is implementable by the treaty X^* if x_i^* maximizes the payoff of player i being on the graph of the restricted best-reply of player j. This result has thus the flavor of the outcome of a sequential game in which player i moves first. Intuitively, this is not surprising since, as already pointed out by Schelling (1960), the power to commit oneself is equivalent to a first move.[13] Hence, implementable profiles

[13]There is now an abundant literature on imperfect competition whose purpose is to obtain Cournot and Stackelberg outcomes as equilibrium outcomes of the same model.

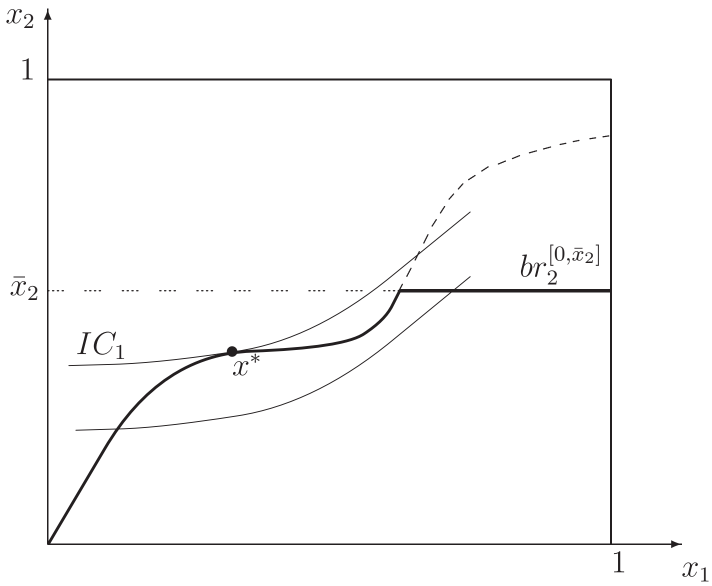

Figure 3: The geometry of Proposition 5

of actions have a Stackelberg-type structure, one player "leads" the treaty while the other "follows."

We now provide a geometric condition that has to hold for a profile of actions to be implementable. In other words, if this condition does not hold at a profile of actions x^* with $x_j^* = BR_j(x_i^*)$, then x^* is *not* implementable; it does not solve the above maximization program. For simplicity, assume that the (mother) best-reply maps and payoff functions are continuously differentiable.[14] The geometric condition relates the slope of the indifference

Interestingly, several models use an approach similar to ours: they give the possibility to the firms to commit to some actions —see for instance Hamilton and Slutsky (1990), van Damme and Hurkens (1999) or more recently Romano and Yildirim (2005), and the references therein. More precisely, firms in most of these models are assumed to commit either to a single action or to not commit at all. Romano and Yildirim (2005) assume that firms can restrict their action sets only from the bottom i.e., firms can only accumulate. Hence these models can be seen as a simplified version of our approach. Hamilton and Slutsky's (1990) main result is that the only equilibrium result that can be obtained are the Cournot and the Stackelberg outcomes, while our approach allows for a larger set of equilibrium results.

[14]The assumption of differentiability is not crucial, but greatly simplifies the exposition.

curve of player i at x^* with the slope of the best-reply of player j at the same action profile x^*.

Proposition 6 *Let x^* be an implementable profile of actions with $x_j^* = BR_j(x_i^*)$, and x^* interior. It cannot be true that the slope of player i's indifference curve at x^* is strictly negative (resp., positive) while the slope of player j's (mother) best-reply at x^* is positive (resp., negative).*

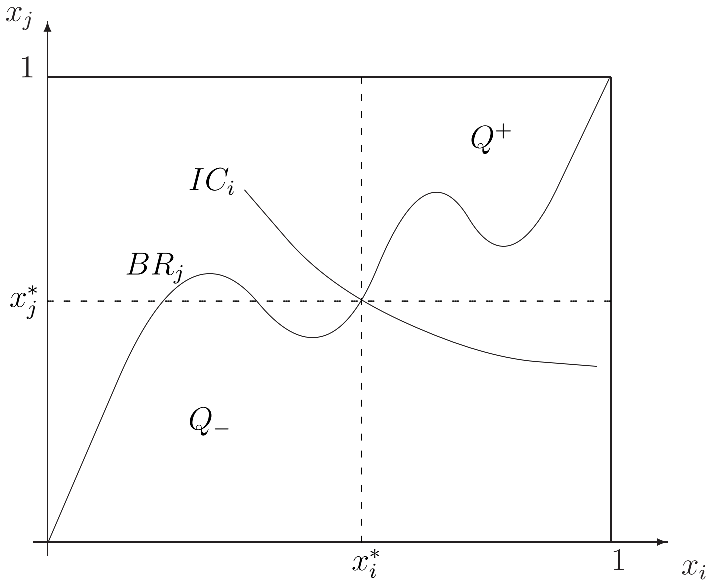

Figure 4: The profile x^* is not implementable.

Proposition 6 thus provides a general geometric condition for implementability: the slope of player i's indifference curve and the slope of player j's best-reply must have the same sign. For instance, in Figure 4, x^* is not implementable since BR_j is positively sloped at x^* while player i's indifference curve IC_i is negatively sloped. Hence, to look for implementable action profiles, we can restrict our attention to the profiles that are on the positively (resp., negatively) sloped portions of the best-reply curve of player j in the positive (resp., negative) indifference curve section of player i. This condition is not sufficient, however. In what follows, we give a necessary and sufficient geometric condition for implementation in an important class of mother games.

Consider the class of games with *strategic complementarities*.[15] Furthermore, we assume that the function $u_i(\cdot, BR_j(\cdot))$ is *strictly quasi-concave* in x_i, for all $i \in \{1,2\}$.[16] For simplicity, we also assume that player i's payoff is increasing in player j's action x_j for all $i \in \{1,2\}$, that is, the game has positive consonance.[17] We show that for this class of games, the knowledge of the Nash equilibria of G along with the knowledge of the "lead-follow" profiles is necessary and sufficient to completely characterize the set of implementable profiles of actions.

First, we need to order the set of Nash equilibria of G. Define $x^*(1)$ the Nash equilibrium of G with the lowest coordinate for player i, that is, there does not exist another equilibrium x of G such that $x_i < x_i^*(1)$. Similarly, define $x^*(2)$ the equilibrium of G with the second lowest coordinate for player i, and so on recursively.[18] Note that since best-reply maps are single-valued, $x^*(k)$ is a singleton for any $k > 0$. Moreover, the set of equilibria of G is generically finite and odd (see Harsanyi (1973)), hence there generically exists a finite odd number K of $x^*(k)$'s. (See Figure 5.)

Second, define $(l_i, BR_j(l_i))$ the profile of actions such that l_i maximizes $u_i(\cdot, BR_j(\cdot))$, that is, the profile of actions $(l_i, BR_j(l_i))$ is the *lead-follow* profile with player i as the leader. It is worth noting that since $u_i(\cdot, BR_i(\cdot))$ is strictly quasi-concave in x_i and BR_j single-valued, l_i is unique. Moreover, since BR_i and u_i are non-decreasing functions of x_j, we have that $l_i \geq x_i^*(K)$ for all $i \in \{1,2\}$ (See Appendix). Our next proposition states that the knowledge of l_i and the $x^*(k)$'s is necessary and sufficient to completely characterize the set of implementable profiles of actions.

[15]See Fudenberg and Tirole (1991, p.490) for a definition. It is worth noting that a similar characterization holds for games with strategic substituabilities.

[16]See Romano and Yildirim (2005) for similar assumptions.

[17]This assumption is not crucial. A complete characterization without this assumption is available upon request.

[18]Formally, let $x^*(0) = \emptyset$, and define for any $k > 0$,

$$x^*(k) := \{x \in N(G) \setminus \cup_{k'=0}^{k-1}\{x^*(k')\} : x_i \leq x'_i,\ \forall x' \in N(G) \setminus \cup_{k'=0}^{k-1}\{x^*(k')\}\}.$$

Before stating the proposition, let us introduce a last piece of notation. Define I_i as a subset of $[0,1]$ as follows:

$$I_i := \bigcup_{\substack{k<K \\ k \text{ odd}}} [x_i^*(k), x_i^*(k+1)] \cup [x_i^*(K), l_i]. \tag{4}$$

Observe that the set I_i is uniquely defined by the knowledge of l_i and the $x^*(k)$'s.

Proposition 7 *Consider a game with strategic complementarities and positive consonance. The set of implementable profiles of actions is $\mathcal{I} = \mathcal{I}_1 \cup \mathcal{I}_2$ with for $i \in \{1,2\}$, $j \neq i$:*

$$\mathcal{I}_i = \{x : x_j = BR_j(x_i),\ x_i \in I_i\}.$$

The intuition behind Proposition 7 is rather simple. First, note that since G is a game with strategic complementarities, the best-reply maps are increasing. Moreover, the best-reply map of any player, BR_i, separates the action space $[0,1]^2$ into two regions $\{x : x_i < BR_i(x_j)\}$ where player i's indifference curves are negatively sloped, and $\{x : x_i > BR_j(x_j)\}$ where player i's indifference curves are positively sloped. Second, for any x with $x_j = BR_j(x_i)$ and $x_i \in (x_i^*(k), x_i^*(k+1))$, k even, we have $x_i < BR_i(x_j)$, hence player i's indifference curve is negatively sloped at x. Since BR_j is positively sloped, it follows from Proposition 6 that x is *not* implementable. A similar argument holds for any x with $x_j = BR_j(x_i)$ and $x_i < x_i^*(1)$. Finally, it is easy to see that any profile of actions x with $x_j = BR_j(x_i)$ and $x_i \in (x_i^*(k), x_i^*(k+1))$, k odd, is implementable by the simple treaty $(\{x_i\}, [0, BR_j(x_i)])$. To see this, it is enough to observe that player j's best-reply $br_j^{[0,BR_j(x_i)]}(x_i')$ is $BR_j(x_j)$ for $x_i' > x_i$, and $BR_j(x_i)$, otherwise. The strict quasi-concavity of u_i and $u_i(\cdot, BR_j(\cdot))$ implies then that x_i is solution of the optimization program described in Proposition 5. The other cases are similar. See Figure 5 for the set of implementable actions.

For the class of games with monotonic best-reply maps and $u_i(\cdot, BR_j(\cdot))$ strictly quasi-concave in x_i, the complete characterization of the set of implementable actions is therefore purely geometric, and the only knowledge required is that of the Nash equilibria of G and the lead-follow profiles.

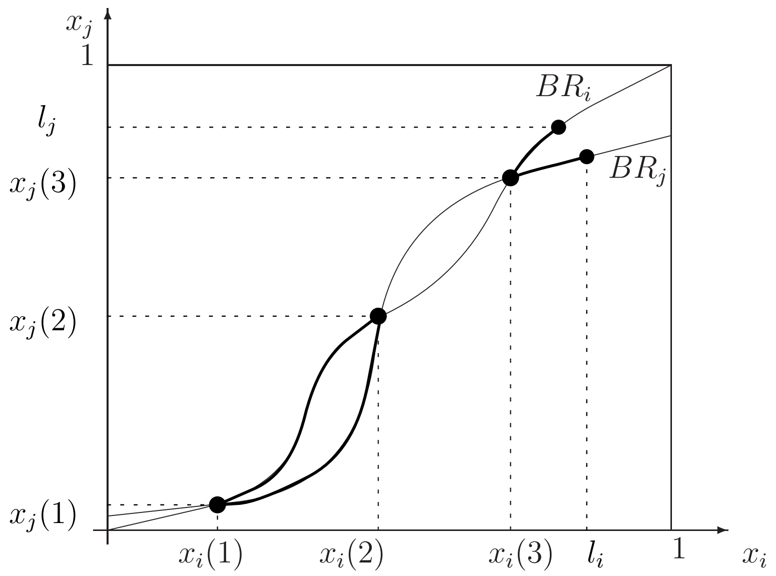

Figure 5: The set of implementable profiles (in bold)

5.5 Examples

We now illustrate our findings by a complete characterization of the set of implementable profiles for a differentiated Bertrand duopoly and a Cournot duopoly. For simplicity, we shall consider the (clasical) model with linear demand, identical cost and such that for both firms the profit function $u_i(x_i, BR_j(x_i))$ is strictly quasi-concave in x_i. Figure 6 and 7 present the best-replies of two firms, 1 and 2, in a differentiated Bertrand duopoly and a Cournot duopoly, respectively. Since firms are symmetric, we shall consider only the profiles x such that $x_2 = BR_2(x_1)$, the profiles where $x_1 = BR_1(x_2)$ are characterized similarly.

In Figure 6, it is easy to see that the strict quasi-concavity of the payoff function implies that all profiles that are in the segment $[A, B]$ are such that firm 1's indifference curve is downward slopping. Thus, using Proposition 6 we deduce that these profiles are not implementable. In Figure 7, the same occurs with the profiles that are in the segment $[l(1), A]$ (with opposite signs) where we define $l(i)$ by $l(i) = (l_i, BR_j(l_i))$.

In Figure 5, profiles beyond $l(1)$ are not implementable either. If they

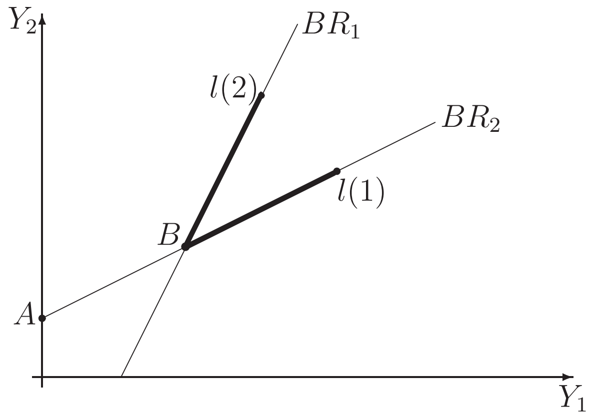

Figure 6: Implementable outcomes in a differentiated Bertrand duopoly.

were implementable, the simple treaty that could be used would either be such that 2's restricted best-reply crosses 1's mother best-reply or such that $l(1)$ is attainable.[19] In Figure 7 those are the points in the segment $[B, C]$.

To complete the characterization of implementable profiles, we can use a modified version of Proposition 7. These profiles are depicted by the bold segments $[B, l(1)] \cup [B, l(2)]$ for the differentiated Bertrand duopoly and $[C, l(1)] \cup [C, l(2)]$ for the Cournot duopoly.

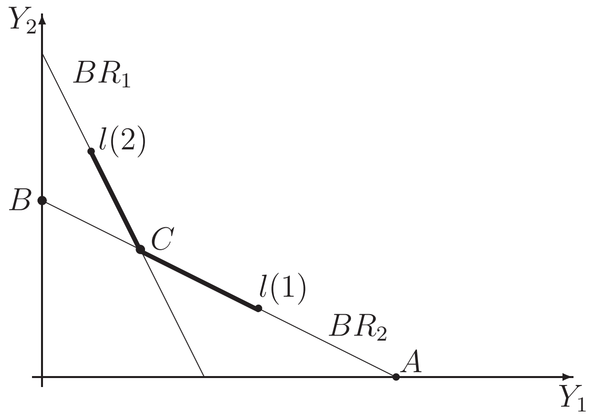

Figure 7: Implementable outcomes in a Cournot duopoly.

[19] Consider x such that $x_1 > l_1$. We only need to scrutinize two different restrictions for firm 2: $[BR_2(x_1), 1]$ and $[0, BR_2(x_1)]$. In the first case firm 1 can do better by choosing the restriction $\{BR_1(BR_2(x_1))\}$, and in the former case by choosing the restriction $\{l_1\}$.

6 The Social Value of Treaties

One conjectured reason for the existence of treaties in real life is that they might make the parties involved in the treaty better off. So it is natural to ask if the treaties in our model make the players better off. In fact, if we interpret the mechanism discussed in this paper as a mechanism to agree on the implementation of particular action profiles we should ask: why don't people just agree to implement efficient profiles? In the case of the example developed in Section 2, we showed that treaties do not give the players the power to implement efficient profiles. In this section, we show that this negative result is not accidental. Treaties quite generally do not implement efficient profiles. More precisely, we show that if G is a smooth game, then we have generic inefficiency.

Next, we address the question whether treaties are at least useful to implement action profiles that Pareto dominate the Nash equilibria of the mother game. We conclude, on a more positive note: we show that treaties can very well serve to make both players better off if certain conditions are met.

6.1 Efficiency

Let us first recall the definition of efficiency.

Definition 4 A profile of actions y is efficient if there does not exist another profile of actions y' such that $u_i(y') \geq u_i(y)$ for all $i \in \{1, 2\}$, and $u_i(y') > u_i(y)$ for some $i \in \{1, 2\}$.

Definition 4 is the textbook definition of (Pareto) efficiency. It is worth noting that several related papers e.g., Jackson and Wilkie (2005) or Gomez and Jehiel (2005), use a stronger concept of efficiency: a profile of actions is efficient if it maximizes the sum of players' payoffs. However, since we assume non-transferable utilities, our concept of efficiency is more appropriate. Let us now turn to the concept of smooth games.

Definition 5 The game G is a smooth game if for all $i \in N$, u_i is twice continuously differentiable.

Two remarks are in order. First, in virtually all economic models in which payoff functions are assumed to be continuous, payoff functions are also assumed to be continuously differentiable.[20] For instance, linear-quadratic Cournot games or models of Bertrand competition with differentiated goods are smooth games. Second, we actually need the assumption of differentiability only around equilibrium results.

THEOREM 2 *For any smooth game G, interior equilibrium results of the treaty game $\Gamma(G)$ are generically inefficient.*[21]

This result is reminiscent of Theorem 1 of Dubey (1986), which states that Nash equilibria of smooth games are generically inefficient. The main reason for hope that this result could be overcome in the game of treaties is that the set of action profiles that can be implemented in this game is (in general a large) superset of the set of Nash equilibria in the same game. So, there is hope that this superset would also contain some efficient profiles. However, our Theorem 2 shows that this does not hold true, just like Nash equilibria of smooth games, the profiles that are implementable by treaties are generically inefficient.

Not only is our Theorem 2 reminiscent of Dubey (1986), also the proof follows along similar lines. The main difference (and difficulty) we face is that implementable profiles that are not themselves Nash equilibria of the mother game lie on the boundary of the action space of the subgame $G(X)$ with X the treaty that is implementing the profile (Lemma 2). This implies that differentiability of the restricted best response fails generically precisely where we need it: at the action profile under investigation.

Some additional remarks are in order. First, our result rests upon the assumption of non-transferable utilities. Allowing for transfers but not for commitment, Jackson and Wilkie (2005) show that efficiency might not hold for two-player games. Whether efficiency holds if we allow for transfers and

[20]Moreover, any continuous function can be arbitrarily approximated by continuously differentiable functions by Weierstrass Approximation Theorem —See Zeidler (1986, p. 770).

[21]Let T be a set of parameters indexing the payoff functions i.e., for each player $i \in \{1, 2\}$, $u_i : X \times T \rightarrow \mathbb{R}$. By genericity, we mean that there exists an open, dense subset of T for which any equilibrium result is inefficient.

commitment is an open question. Second, Theorem 2 continues to hold if G is a game with strategic complementarities, but not necessarily smooth. (See Appendix.) Third, efficient profiles on the boundary can in some games be implemented by treaties. This holds in particular if a game has an efficient Nash equilibrium on the boundary.

6.2 Pareto Improvements

While efficient results are generically not implementable, a self-enforcing treaty might nonetheless implement an improvement upon the status quo. In other words, the next question we address is whether a treaty can implement a profile that makes both players better off compared to any equilibrium of the mother game G.

Definition 6 A result x^* is an *improvement upon the status quo* if $u_i(x^*) \geq u_i(y^*)$ for all $i \in \{1, 2\}$, and $u_i(x^*) > u_i(y)$ for at least one player, where y^* is an action profile that is efficient in the set of mother Nash equilibria.[22]

It is not hard to find games in which improvements upon the status quo can be implemented. Just take any game with a unique Nash equilibrium y^* and a follow-lead-equilibrium that dominates y^*.[23] The follow-lead equilibrium can be implemented by the treaty in which the leader restricts his action space to a singleton while the follower does not restrict his action space at all. So the more interesting question is: can treaties be used to implement improvements upon the status quo if none of the follow-lead equilibria represents such an improvement? In our next result we show that this cannot happen if the players' best responses are monotone and if the players' utilities are monotone in the actions of the opponents. We say that a game satisfies *constant consonance* if any players payoff is monotone in the action of the other player.

[22]Note that the set of equilibria $N(G)$ is a compact set, hence efficiency is well defined.

[23]This is the case for instance of all symmetric games with a second-mover advantage. Since the payoff of the first player in a follow-lead profile is necessarily weakly higher than the highest Nash equilibrium, the former Pareto dominates the latter.

Theorem 3 *Let G be a game with constant consonance such that the follow-lead equilibria do not improve on the status quo. Then there exists an equilibrium improvement x^* only if at least one best-reply map is non-monotonic.*

An important implication of Theorem 3 is that if G is a game with strategic complementarities or strategic substituabilties, then treaties do only serve to improve upon the status quo if the follow-lead equilibrium is already itself such an improvement. This result sharply contrasts with Proposition 2 of Bernheim and Whinston (1998), and illustrates how seemingly innocuous restrictions on the set of feasible commitments can be critical. Bernheim and Whinston's model and our model, albeit similar in spirit, differ in two important dimensions. First, Bernheim and Whinston allow for non-convex restrictions while we do not. Second in their model only a one player (the principal) has the opportunity to commit. Thirdly, and more importantly, this player does not only have the power to commit himself (to take a single action) he can also restrict the action set of the other player, the agent. In contrast, both players have the power to commit in our model.

Theorems 2 and 3 are rather negative results in that the power of commitment does not seem to be of much social value. The following example shows that equilibrium improvements do exist even in the case that neither of the follow-lead equilibria represents such an improvement.[24]

Example 1 Take the mother game G with strategy spaces $Y_1 = Y_2 = [0, 2]$ and payoff functions:

$$u_1(y_1, y_2) = \frac{y_1}{\frac{y_1}{4} + y_2} - y_1,$$

$$u_2(y_1, y_2) = -\left(y_2 + \frac{y_1}{2} - \frac{2}{3}\right)^2 - 20y_1.$$

The best-reply map of the players are

$$BR_1(y_2) = \begin{cases} -4y_2 + 4\sqrt{y_2} & \text{if } y_2 \leq 1, \\ 0 & \text{otherwise,} \end{cases}$$

[24] The present example differs from the example in section 2 insofar as that the follow-lead equilibria in the example of section 2 Pareto dominates the unique Nash equilibrium of that game.

and

$$BR_2(y_1) = \begin{cases} -\dfrac{1}{2}y_1 + \dfrac{2}{3} & \text{if } y_1 \leq \dfrac{4}{3}, \\ 0 & \text{otherwise.} \end{cases}$$

The mother game has a unique equilibrium, $y_1^* = 4/3(\sqrt{3}-1)$, $y_2^* = 2/3(2-\sqrt{3})$, with equilibrium payoffs of $u_i(y^*) = 4/3$, $u_j(y^*) = 80/3(1-\sqrt{3}) \simeq -19.52$, respectively. Moreover, the follow-lead profile $(BR_1(l_2), l_2) = (1, 0)$ is associated to payoffs of $u_1((BR_1(l_2), l_2))) = 0$, $u_2((BR_1(l_2), l_2)) = -1/9 \simeq -0.11$.

Let us show that there exists a self-enforcing treaty which implements the action profile $\tilde{y} = (8/9, 1/9)$ with associated payoffs of $u_1(\tilde{y}) = 16/9$ and $u_2(\tilde{y}) = -1441/81 \simeq -17.79$, respectively. Clearly, both players' payoffs improve upon the Nash equilibrium. According to Proposition 2, at least one player's action must be a best-reply against the action of the other player. In the profile $\tilde{y}$, we have $8/9 = BR_1(1/9)$.

Following Proposition 4, we can focus, without loss of generality, on only two candidates for the restriction of player 1, $[0, 8/9]$ or $[8/9, 1]$. We claim that player 1's restriction cannot be $[0, 8/9]$. To see this, observe that if 1 commits to $[0, 8/9]$, then player 2 can commit to $\{1\}$ and gets a payoff of $-1/9$ (since $br_1^{[0,8/9]}(1) = 0$), which is higher than $u_2(\tilde{y})$. Therefore, the unique candidate for 1's restriction is $[8/9, 1]$. In this case, player 1's restricted best-reply is

$$br_1(y_2) = \max\{-4y_2 + 4\sqrt{y_2}, 8/9\}. \tag{5}$$

Observe that for all $y_2 \in [1/9, 4/9]$, we have $-4y_2 + 4\sqrt{y_2} \geq 8/9$. It follows that 2's payoff when $y_2 \notin [1/9, 4/9]$ is $-(y_2 - 2/9)^2 - 160/9$, which is maximized when $y_2 = 1/9$. If $y_2 \in [1/9, 4/9]$, then player 2 maximizes $u_2(y) = -4y_2 + 4\sqrt{y_2}$. That the maximum is obtained when $y_2 = 8/9$ is a simple matter of computation (albeit tedious) and is left to the reader.

Appendix

Proof of Proposition 2 The proof proceeds by contradiction. Let $s^* = (X_i^*, \sigma_i^*)_{i\in\{1,2\}}$ be an equilibrium of Γ, and suppose that (X^*, x^*) the outcome of s^* is such that $x_i^* \neq BR_i(x_j^*)$, $i \neq j$, $i \in \{1,2\}$. To reach a contradiction, we will first identify an action, x_1' such that $u_1(x_1', x_2^*) > u_1(x_1^*, x_2^*)$, and second, show that there exists a strategy for player 1 s_1' such that the outcome of (s_1', s_2^*) is $(x^*, (x_1', x_2^*))$, hence a contradiction with s^* being an equilibrium.

Step 1. Since x^* is a Nash equilibrium of the game $G(X^*)$, we have $x_i^* = br_i^{X^*}(x_j^*)$ for $i \in \{1,2\}$, $i \neq j$. By our supposition we have that $br_i^{X^*}(x_j^*) \neq BR_i(x_j^*)$ for $i \in \{1,2\}$, $i \neq j$. Hence there exists an open ball $\mathcal{B}_\varepsilon(x_1^*)$ of center x_1^* and sufficiently small radius ε such that for all $x_1 \in \mathcal{B}_\varepsilon(x_1^*)$ we have that $br_2^{X^*}(x_1) = x_2^*$. Next pick $\alpha \in [0,1)$ large enough enough such that $x_1' = \alpha x_1^* + (1-\alpha)BR_1(x_2^*) \in \mathcal{B}_\varepsilon(x_1^*)$. By the construction of $\mathcal{B}_\varepsilon(x_1^*)$ we have that $br_2^{X^*}(x_1') = x_2^*$, and due to strict quasi-concavity of player 1's preferences we have that $u_1(x_1', x_2^*) > u_1(x_1^*, x_2^*)$.

Step 2 We claim that the strategy $s_1' = (\{x_1'\}, \sigma_1^*)$ is a profitable deviation for player 1. The outcome of (s_1', s_2^*) is $((\{x_1'\}, X_2^*), (x_1', x_2^*))$, which player 1 by construction prefers to x^*. ∎

Proof of Proposition 3 Let $s^* = ((X_1^*, \sigma_1^*), (X_2^*, \sigma_2^*))$ be an equilibrium of Γ with outcome (X^*, x^*). By Proposition 2, for one player, say 1, we have $x_1^* = BR_1(x_2^*)$. We claim that strategy profile $s' := (s_1^*, s_2')$, with $s_2' = (\{x_2^*\}, \sigma_2^*)$ is also an equilibrium, we furthermore claim that $((X_1^*, \{x_2^*\}), x^*)$ is the outcome of this strategy profile. The second claim is easy to see as $x_1^* = BR_1(x_2^*) \in X_1^*$.

For the first claim observe first that player 1 does not have an incentive to deviate: Since 2's restriction is the singleton $\{x_2^*\}$, 1 cannot possibly obtain more than $u_1(BR_1(x_2^*), x_2^*)$, which is exactly what he obtains for s'. Hence, there is no profitable deviation for player 1. To show that player 2 has no profitable deviation either, we use the one shot deviation property. In the second stage player 2's strategy is prescribed by σ_2^* which by our assumption that s^* is an equilibrium yields equilibrium play in all proper subgames.

Consequently player 2 has no profitable deviations for second stage play. Suppose now that $s_2'' = (X_2'', \sigma_2^*)$ where a profitable deviation for player 2, i.e player 2 gets a higher utility from (s_1^*, s_2'') than from (s_1^*, s_2'). This cannot be: By construction player 2 is indifferent between (s_1^*, s_2') and s^*, so also for s^* would s_2'' be a profitable deviation for player 2. This stands in contradiction with our assumption that s^* is an equilibrium. ∎

Proof of Proposition 4 Let $s^* = ((\{x_i^*\}, \sigma_i^*), (X_j^*, \sigma_j^*))$ be an equilibrium of Γ with result x^*. Define $s_j' = ([x_j^*, 1], \sigma_j^*)$ and $s_j'' = ([0, x_j^*], \sigma_j^*)$. We claim that either s_i^*, s_j' or s_i^*, s_j'' is an equilibrium of Γ with result x^*. First observe that both of the strategy profiles under consideration have x^* as their result. This is so since in either case player i is restricting their action space to $\{x_i^*\}$ and in either case x_j^* player j's mother best-reply to x_i^* is part of the restriction chosen by player j. Next observe that for s_i^*, s_j' or s_i^*, s_j'' we have that x^* is an equilibrium of the restricted games $G(\{x_i^*\}, [x_j^*, 1])$ and $G(\{x_i^*\}, [0, x_j^*])$. In both games, player i has only one action, and player j's best response to that action according to the mother game $x_j^* = BR_j(x_i^*)$ is contained in the restricted action space in either case. It is equally easy to see that player j does not have an incentive to change his restricted action space as his best response to the only action in player i's restricted action space is contained in the new restricted action space of player j.

Next we will establish that for player i the choice of the action space $\{x_i^*\}$ is either a best response when player j chooses action space $[x_j^*, 1]$ or the alternative action space $[0, x_j^*]$. Since s^* is an equilibrium, the set of outcomes that player i strictly prefers to x^* lies either above or below the restricted best response of player j. If not this set could not be convex, as no action profile on the restricted best response curve of player j can belong to this set. So we have either that:

$$br_j^{[\underline{x}_j, \overline{x}_j]}(x_i') - x_j' > 0 \text{ for all } x' \text{ such that } u_i(x') > u_i(x^*)$$

or

$$br_j^{[\underline{x}_2, \overline{x}_2]}(x_i') - x_j' < 0 \text{ for all } x' \text{ such that } u_i(x') > u_i(x^*)$$

Next observe that

$$\begin{array}{r} br_2^{[\underline{x}_2,\overline{x}_2]}(x_1) \leq br_2^{[x_2^*,\overline{x}_2]}(x_1) \leq br_2^{[x_2^*,1]}(x_1) \\ br_2^{[\underline{x}_2,\overline{x}_2]}(x_1) \geq br_2^{[\underline{x}_2,x_2^*]}(x_1) \geq br_2^{[0,x_2^*]}(x_1) \\ \text{for all } x_1 \in Y_1 \end{array}$$

So in the first case we obtain that:

$$br_2^{[x_2^*,1]}(x_1') - x_2' > 0 \text{ for all } x' \text{ such that } u_1(x') > u_1(x^*)$$

This implies that given that player 2 picks action space $[x_2^*, 1]$ player 1 cannot obtain a utility higher than $u(x^*)$. In other words, the choice of the action space $\{x_1^*\}$ is optimal for player 1.

If the second case holds true we obtain that

$$br_2^{[0,x_2^*]}(x_1') - x_2' > 0 \text{ for all } x' \text{ such that } u_1(x') > u_1(x^*)$$

Which would of course imply that player 1 cannot do better than picking action space $\{x_1^*\}$ when his opponent picks $[0, x_2^*]$.

We conclude that either in s' or in s'' both players are best responding. So either one must be a an equilibrium on Γ. ∎

Proof of Proposition 5 Observe that we can rewrite conditions *(i)* and *(ii)* as follows. A profile (x_1^*, x_2^*) is implementable by a treaty if and only if there exists a restriction X_j^* such that x_i^* is a solution of the following program,

$$\begin{cases} (\mathcal{P}) \begin{cases} \max_{x_i} u_i(x_i, x_j) \\ \text{s.t. } (x_i, x_j) \in \mathrm{Gr}(br_j^{X_j^*}) \end{cases} \\ \text{such that } br_j^{X_j^*}(x_i^*) = BR_j(x_i^*). \end{cases} \qquad (\mathcal{P}^*)$$

where $i, j = 1, 2$ and $i \neq j$. Note that $(\mathcal{P}^*)$ is a two-step optimisation program. First, we optimise $u_i(x_i, br_j^{X_j^*}(x_i))$ with respect to x_i. This is the program $(\mathcal{P})$. Second, we check whether the solution obtained lies on the graph of j's best-reply BR_j.

$(\Rightarrow)$ Let $s^* = (X_i, \sigma_i^*)_{i=1,2}$ be an equilibrium of Γ, where $X_1^* = \{x_1^*\}$. (The case when $X_2^* = \{x_2^*\}$ is symmetric). For all $X \in \mathcal{Y}$, the mappings

σ_1 and σ_2 are such that $(\sigma_1^*(X), \sigma_2^*(X))$ is a Nash equilibrium of $G(X)$. In particular, if $X_1 = \{x_1\}$ for some $x_1 \in Y_1$, we have $\sigma_2^*(X) = br_2^{X_2}(x_1)$. Thus, for all deviations by player 1 to a strategy $s_1 = (\{x_1\}, \sigma_1^*)$ for some $x_1 \in Y_1$, we have $u_1(s_1, s_2^*) = u_1(x_1, br_2^{X_2^*}(x_1))$. Since s^* is an equilibrium, such deviations are not profitable, i.e.,

$$u_1(x_1^*, br_2^{X_2^*}(x_1^*)) \geq u_1(x_1, br_2^{X_2^*}(x_1)), \forall\, x_1 \in Y_1\,.$$

That is, x_1^* must be a solution of $(\mathcal{P})$. By Proposition 2, we have $x_i^* = BR_i(x_{-i}^*)$ for at least one player $i \in \{1, 2\}$. Suppose that $x_2^* \neq BR_2(x_1^*)$. Then, given $(\{x_1^*\}, x_1^*)$, player 2 is better-off deviating to $(X_2', BR_2(x_1^*))$, a contradiction. Hence, we have $x_2^* = BR_2(x_1^*)$. Thus, x_1^* is solution of $(\mathcal{P}^*)$.

$(\Leftarrow)$ Suppose that x_1^* is solution of $(\mathcal{P}^*)$. Consider the following strategy profile: $s_1^* = (\{x_1^*\}, \sigma_1^*)$, and $s_2^* = (X_2^*, \sigma_2^*)$, where the mappings σ_1^* and σ_2^* are such that $(\sigma_1^*(X), \sigma_2^*(X))$ is a Nash equilibrium of $G(X)$, for all $X \in \mathcal{Y}$. Clearly, the outcome of s^* is (x_1^*, x_2^*), and by construction it also is a Nash equilibrium of $G(\{x_1^*\} \times X_2^*)$.[25] By construction, for all subgames $G(X)$, the actions $(\sigma_1^*(X), \sigma_2^*(X))$ constitute a Nash equilibrium of $G(X)$. Hence, according to the one-shot deviation property, it suffices to check that there is no first-stage deviation to obtain that s^* is indeed an equilibrium of Γ. Since $x_2^* = BR_2(x_1^*)$ and $X_1^* = \{x_1^*\}$, player 2 cannot obtain a better payoff than $u_2(x_1^*, x_2^*)$, and thus does not deviate. As for player 1, suppose that $\exists\, X_1 \in \mathcal{Y}_1$ such that for $s_1 = (X_1, \sigma_1^*)$, $u_1(s_1, s_2^*) > u_1(s_1^*, s_2^*)$. Let (x_1, x_2^*) be the outcome of the profile (s_1, s_2^*). Since s_1 is a profitable deviation, we then have $u_1(x_1, x_2^*) > u_1(x_1^*, x_2^*)$. By construction (of the mapping σ_2), $x_2^* = br_2^{X_2^*}(x_1)$, a contradiction with the fact that x_1^* is a solution of $(\mathcal{P})$. ∎

Proof related to the geometry. We first start with a preliminary result.

Lemma 3 *Let G be a game with strategic complementarities and positive*

[25] Since x_1^* is solution of $(\mathcal{P}^*)$, $br_2^{X_2^*}(x_1^*) = BR_2(x_1^*) \in X_2^*$. Moreover, single-valuedness of BR_2 implies that (x_1^*, x_2^*) is the unique Nash equilibrium of $G(\{x_1^*\} \times X_2^*)$, where $x_2^* = BR_2(x_1^*)$.

consonance i.e., u_i is non-decreasing in x_j, $j \neq i$, for all $i \in N$. We have $l_i \geq x_i^*(K)$.

Proof Suppose that $x_i^*(k+1) > l_i > x_i^*(k)$. Since, BR_j is non-decreasing, we have $BR_j(x_i^*(k+1)) \geq BR_j(l_i) \geq BR_j(x_i^*(k))$, hence

$$u_i(l_i, BR_j(x_i^*(k+1)) \geq u_i(l_i, BR_j(l_i)) \tag{7}$$

since u_i has positive consonance. Moreover, since $x_i^*(k+1)$ is the unique best-reply to $x_j^*(k+1) = BR_j(x_i^*(k+1))$ ($x^*(k+1)$ is a Nash equilibrium), we have

$$\begin{aligned} u_i(x_i^*(k+1), x_j(k+1)) &> u_i(l_i, BR_j(x_i^*(k+1)) \\ &\geq u_i(l_i, BR_j(l_i)) \geq u_i(x_i^*(k+1), x_j^*(k+1)), \end{aligned} \tag{8}$$

a contradiction. A similar argument shows that l_i could not be smaller than $x_i^*(1)$. ■

Proof of Proposition 6 Let x^* be an implementable profile of actions with $x_j^* = BR_j(x_i^*)$, and x^* interior. By contradiction, suppose that the slope of indifference curve of player i at x^* is negative while the slope of BR_j at x^* is positive.

Define $Q^+ := \{y \in [0,1]^2 : y \geq x^*\}$ and $Q_- = \{y \in [0,1]^2 : y \leq x^*\}$.[26] Since the indifference curve of player i at x^* is negatively sloped, there exists either an ε_1 such that $u_i(y) > u_i(x^*)$ for all $y \in B_{\varepsilon_1}(x^*) \cap (Q^+ \setminus \{x^*\})$ or an ε_2 such that $u_i(y) > u_i(x^*)$ for all $y \in B_{\varepsilon_2}(x^*) \cap (Q^- \setminus \{x^*\})$. Let $\varepsilon := \min(\varepsilon_1, \varepsilon_2)$.

Let $f : X \to Y$ be a function. We denote $\mathrm{Gr}\, f$ the graph of f. Since the slope of BR_j at x^* is positive, we have that

$$\begin{aligned} &\mathrm{Gr}\; br_j^{[0, BR_j(x_i^*)]} \cap (B_\varepsilon(x^*) \cap Q^+ \setminus \{x^*\}), \\ &\mathrm{Gr}\; br_j^{[0, BR_j(x_i^*)]} \cap (B_\varepsilon(x^*) \cap Q_- \setminus \{x^*\}), \\ &\mathrm{Gr}\; br_j^{[BR_j(x_i^*), 1]} \cap (B_\varepsilon(x^*) \cap Q^+ \setminus \{x^*\}), \\ &\mathrm{Gr}\; br_j^{[BR_j(x_i^*), 1]} \cap (B_\varepsilon(x^*) \cap Q_- \setminus \{x^*\}), \end{aligned}$$

[26] Let x and y two vectors in $\mathbb{R}^n$. We write $x \geq y$ if $x_i \geq y_i$ for all $i \in \{1, \ldots, n\}$

are non-empty sets.

Finally, from Theorem 1, the two simple treaties that could possibly implement the profile x^* are $(\{x_i^*\}, [0, BR_j(x_i^*)])$ and $(\{x_i^*\}, [BR_j(x_i^*), 1])$. It follows from the above arguments that x^* cannot be a solution of the optimization program described in Proposition 5, hence a contradiction with x^* being implementable. The same argument follows mutatis mutandum for the other cases. ■

Proof of Proposition 7 We first start with a preliminary observation. The best-reply of player i separates the action space $[0,1]^2$ into two regions: one region in which player i's indifference curves are negatively sloped, one region in which player i's indifference curves are positively sloped. To prove this result, fix an action x_j^* of player j, and consider the best-reply $x_i^* = BR_i(x_j^*)$ of player i to x_j^*. Define $IC := \{x \in [0,1]^2 : u_i(x) = u_i(x^*)\}$. For any $x_i \neq x_i^*$, we have $u_i(x_i, x_j^*) < u_i(x^*)$ since x_i^* is the unique best-reply to x_j^*. Next, if $x_j < x_j^*$, it follows from u_i increasing in x_j that $u_i(x_i, x_j) \leq u_i(x_i, x_j^*) < u_i(x^*)$, hence $(x_i, x_j) \notin IC$. Therefore, for any x_i, we need $x_j > x_j^*$ for (x_i, x_j) to belong to IC. Hence, we have that for any $x_i < x_i^*$, IC is negatively sloped and for any $x_i > x_i^*$, IC is positively sloped.

As a second observation, observe that for any $x_i \in [x_i^*(k), x_i^*(k+1)]$, $BR_i(BR_j(x_i)) - x_i$ is either positive or negative, but does not alternate in signs. For otherwise, there exists another equilibrium in $[x_i^*(k), x_i^*(k+1)]$, a contradiction with the definition of the $x^*(k)$'s. Moreover, we have that $BR_i(BR_j(x_i)) - x_i \leq 0$ for any $x_i \in [x_i^*(k), x_i^*(k+1)]$ if k is odd, $BR_i(BR_j(x_i)) - x_i \geq 0$, if k is even. In words, the graph of player i's best-reply is to the "left" of the graph of player j's best-reply if k is odd, and to the "right" if k is even. (See Figure 5.) Furthermore, $BR_i(BR_j(x_i)) - x_i > 0$ for any $x_i < x_i^*(1)$ and $BR_i(BR_j(x_i)) - x_i < 0$ for any $x_i > x_i^*(K)$.[27]

Fix a profile of actions x with $x_j = BR_j(x_i)$ and $x_i \in (x_i^*(k), x_i^*(k+1))$ for some k even. We want to show that this profile is not implementable.

[27]By contradiction, suppose that $BR_i(BR_j(x_i)) - x_i < 0$ for any $x_i < x_i(1)$. In particular, for $x_i = 0$, i.e., for the lower bound of Y_i, we have $0 \leq BR_i(BR_j(0)) - 0 < 0$, a contradiction.

From the previous observation, we have that $BR_i(x_j) = BR_i(BR_j(x_i)) > x_i$. From the first observation, it then follows that the indifference curve of player i at x is negatively sloped. Since BR_j is positively sloped, it follows from Proposition 6 that x is not implementable. A similar argument holds for any x with $x_j = BR_j(x_i)$ and $x_i < x_i^*(1)$.

Let us now consider any profile of actions x^* with $x_j^* = BR_j(x_i^*)$ and $x_i^* \in (x_i^*(k), x_i^*(k+1))$ for some k odd. We want to show that any such a profile is implementable by the simple treaty $(\{x_i^*\}, [0, BR_j(x_i^*)])$. The key observation is that the best-reply of player i is now to the "left" of the best-reply of player j i.e., $BR_i(BR_j(x_i^*)) < x_i^*$. (See Figure 5.) Hence, for any $x_i > x_i^*$, $br_j^{X_j^*}(x_i) = BR_j(x_i^*)$, that is, player j's restricted best-reply is $BR_j(x_i^*)$, and $u_i(x_i, br_j^{X_j^*}(x_i)) < u_i(x_i^*, br_j^{X_j^*}(x_i^*))$ by strict quasi-concavity of u_i. Finally, note that $br_j^{X_j^*}(x_i) = BR_j(x_i)$ for any $x_i \leq x_i^*$, henceforth the maximum of $u_i(\cdot, br_j^{X_j^*}(\cdot))$ is achieved in x_i^* by strict quasi-concavity of $u_i(\cdot, BR_j(\cdot))$. It follows that x^* is implementable (step 4).

Similar arguments complete the proof. For instance, any point x^* with $x_j^* = BR_j(x_i^*)$ and $x_i^* \in (x_i^*(K), l_i]$ is implementable by the simple treaty $(\{x_i^*\}, [0, BR_j(x_i^*)])$. ■

Proof of Theorem 2 Let (x^*, X^*) be any equilibrium outcome of $\Gamma(G)$ such that x^* is interior.

Let T be a set of parameters and define the family of payoff functions : $u_i : X \times T \to \mathbb{R}$, for all $i \in \{1, 2\}$. We want to show that for a dense open subset T^* of T, x^* is inefficient.

If x^* is an equilibrium of the mother game G, the result follows from Theorem 1 of Dubey (1986).

If x^* is not an equilibrium of the mother game G, the proof is similar to the proof of Theorem 1 of Dubey. The proof is as follows.

Define the directional mapping $D : T \times X \to \mathbb{R}^4$ with

$$D(t, x') = \begin{bmatrix} \frac{\partial u_1(\cdot,t)}{\partial x_1}(x') & \frac{\partial u_1(\cdot,t)}{\partial x_2}(x') \\ \frac{\partial u_2(\cdot,t)}{\partial x_1}(x') & \frac{\partial u_2(\cdot,t)}{\partial x_2}(x') \end{bmatrix} \tag{9}$$

and let $D_t(\cdot)$ be the restriction of D to t. Thus, $D_t(x^*)$ is the Jacobian matrix evaluated at x^*. The key step in Dubey's proof is to observe that at any interior equilibrium x^* of G, the diagonal elements of the Jacobian matrix are zero, and that the set of 2×2 matrices with zeros on the diagonal is a submanifold of $\mathbb{R}^4$ of co-dimension 2. If x^* is *not* an equilibrium of G, we have a similar result, that is, we can show that if x^* is an equilibrium of Γ, then $D_t(x^*) \in A \cap B$, with $A \cap B$ a submanifold of $\mathbb{R}^4$ of co-dimension 2. This step is the only step that differs with Dubey's proof.

First, from Lemma 2, for at least one player, we have $x_i^* = BR_i(x_{-i})$. Without loss of generality, suppose that $x_2^* = BR_2(x_1^*)$. Since x^* is interior, we then have that $\frac{\partial u_2}{\partial x_2}(x^*) = 0$. This equality is our first constraint on the Jacobian matrix. Formally, define the set

$$A = \{M \in \mathbb{R}^4 : M_{22} = 0\}, \tag{10}$$

i.e., the set of 2×2 matrices with a zero on the diagonal. Observe that if x^* is an equilibrium result, then $D_t(x^*) \in A$, or $x^* \in D_t^{-1}(A)$. The set A is a submanifold of $\mathbb{R}^4$ of co-dimension 1.

Second, since (x^*, X^*) is an equilibrium outcome, it follows from Theorem 1 that $u_1(x_1^*, br_2^{X_2^*}(x_1^*)) \geq u_1(x_1, br_2^{X_2^*}(x_1))$ for all $x_1 \in Y_1$. We show that these inequalities impose a relationship between the first-order derivatives of u_1 with respect to x_1 and x_2, respectively. If $br_2^{X_j^*}$ is differentiable at x^*, then the relationship is similar to the one in the text. However, whenever X^* is a simple treaty, $br_2^{X_2^*}$ is not differentiable in x_1^*. We use the concepts of subgradient and subdifferential to circumvent this problem.[28]

For any function $f : Z \to \mathbb{R}$, denote $\partial f(z)$ the subdifferential of f at z. We refer the reader to Clarke (1989, Chapter 1) or Rockafellar (1981, Chapter 3) for rigorous definitions of subdifferentials. As an example, if $f(z) = |z|$, then $\partial f(0) = [-1, 1]$.

Since u_2 is twice continuously differentiable, BR_2 is continuously differentiable, hence Lipschitz continuous. From Lemma 1, it then follows that

[28]We refer the reader to Rockafellar (1981) for a good source on the theory of subgradients and non-smooth optimisation.

$br_2^{X_2^*}$ is Lipschitz continuous. Note that Rademacher Theorem implies that $br_2^{X_2^*}$ is differentiable almost everywhere. Let us consider the subdiffential of $v_1(\cdot) := -u_1(\cdot, br_2^{X_2^*}(\cdot))$ at x_1^*. Since u_1 is continuously differentiable and $br_2^{X_2^*}$ is Lipschitz continuous, Theorem 5P of Rockafellar (Rockafellar 1981, p 74) implies that

$$\partial v_1(x_1^*) = -\frac{\partial u_1}{\partial x_1}(x^*) - \frac{\partial u_1}{\partial x_2}(x^*)\partial br_2^{X_2^*}(x_1^*). \tag{11}$$

Since x_1^* minimises v_1, $0 \in \partial v_1(x_1^*)$ (Clarke (1989, p. 9)), hence there exists a $\xi \in \partial br_2^{X_2^*}(x_1^*)$ such that

$$0 = \frac{\partial u_1}{\partial x_1}(x^*) + \frac{\partial u_1}{\partial x_2}(x^*)\xi, \tag{12}$$

the required relationship. (Note that if $br_2^{X_2^*}$ is differentiable at x_1^*, then ξ is the derivative of $br_2^{X_2^*}$ evaluated at x_1^*.)

For any scalar a, define the set

$$B = \{M \in \mathbb{R}^4 : M_{11} + aM_{12} = 0\}, \tag{13}$$

i.e., the set of 2×2 matrices with a linear relationship between the two first entries. It follows that if x^* is an equilibrium result, then $D_t(x^*) \in B$, or $x^* \in D_t^{-1}(B)$ (take $a = \xi$). The set B is a submanifold of $\mathbb{R}^4$ of co-dimension 1. It then trivially follows that $A \cap B$ is a submanifold of $\mathbb{R}^4$ of co-dimension 2, as required.

Finally, define the set

$$C = \{M \in \mathbb{R}^4 : \text{the rows of } M \text{ are linearly dependent}\}. \tag{14}$$

It is easy to see that if x^* is efficient, then $D_t(x^*) \in C$, or $x^* \in D_t^{-1}(C)$. For otherwise, there exists a neighbourhood O of x^* and a $x' \in O$ such that $u_i(x') = u_i(x^*) + \varepsilon_i$, $\varepsilon_i > 0$, for all player $i \in N$ i.e., there exists dx_1 and dx_2 such that

$$\begin{bmatrix} \frac{\partial u_1(\cdot,t)}{\partial x_1}(x^*) & \frac{\partial u_1(\cdot,t)}{\partial x_2}(x^*) \\ \frac{\partial u_2(\cdot,t)}{\partial x_1}(x^*) & \frac{\partial u_2(\cdot,t)}{\partial x_2}(x^*) \end{bmatrix} \begin{pmatrix} dx_1 \\ dx_2 \end{pmatrix} = \begin{pmatrix} \varepsilon_1 \\ \varepsilon_2 \end{pmatrix}. \tag{15}$$

Hence, if a profile x^* is an equilibrium result and efficient, then $D_t(x^*) \in A \cap B \cap C$ or $x^* \in D_t^{-1}(A \cap B \cap C)$.

The next step is to show that for a dense open set $T^* \subset T$, $D_t^{-1}(A \cap B \cap C)$ is empty. To do so, we shall show that the co-dimension of $D_t^{-1}(A \cap B \cap C)$ is 2, that is the dimension of Y, hence is empty. This step is found in Dubey's proof. ∎

Inefficiency and a non-smooth game.

Consider the game G to be a game with strategic complementarities and negative consonance i.e., $x_j \mapsto u_i(x_i, x_j)$ is decreasing in x_j for each player $i \in N$, $i \neq j$. Note that G is not assumed to be smooth.

The first observation is that $BR_1(BR_2(x_1^*)) \leq x_1^*$. Since BR_2 is monotone increasing in x_1, we have $br_2^{[0, BR_2(x_1^*)]}(x_1) = BR_2(x_2)$ for all $x_2 \in [0, x_1^*]$, and $br_2^{[0, BR_2(x_1^*)]}(x_1) = BR_2(x_1^*)$, otherwise. Henceforth, if $BR_1(BR_2(x_1^*)) > x_1^*$, we have that player 2's best-reply to $BR_1(BR_2(x_1^*))$ is $BR_2(x_1^*)$, hence a contradiction with x_1^* maximising player 1's payoff on the constrained best-reply of player 2.

Second, since u_2 is decreasing in x_1, we obviously have

$$u_2(x_2^*, BR_1(BR_2(x_1^*))) \geq u_2(x_2^*, x_1^*),$$

hence $(BR_1(BR_2(x_1^*)), x_2^*)$ improves upon 2's payoff.

Finally, since at an equilibrium x^* of Γ, $x_2^* = BR_2(x_1^*)$, it follows that

$$u_1(BR_1(x_2^*), x_2^*) \geq u_1(x^*),$$

with a strict inequality if x^* is not a Nash equilibrium of G.

It follows that $(BR_1(x_2^*), x_2^*)$ Pareto-improves upon x^*, hence x^* is not efficient.

Finally, observe that the result also holds if we assume strategic substitutes and payoff increasing in the action of the opponent.

Proof of Theorem 3

Let (X^*, x^*) be an equilibrium outcome of Γ and assume that x^* is an improvement on the status quo. Let x^N be the Nash equilibrium, that is efficient in the set of Nash equilibria, for which we have that $u_i(x^*) \geq u_i(x^N)$ for $i = 1, 2$ with at least one strict inequality. Using Lemma 1, we can assume that $x_2^* = BR_2(x_1^*)$. By our assumption that neither of the follow-lead equilibria is an improvement upon the status quo we have that

$$u_2(x^*) \geq u_2(x^N) > u_2(l_1, BR_2(l_1)).$$

Observe that in all the three profiles player 2 is best responding. Furthermore, as player 2's payoff function is monotonic in his opponent's action, we have that $u_2^*(x_1) := u_2(x_1, BR_2(x_1))$ is a monotonic function of x_1, hence x_1^* and l_1 must lie on two different sides of x_1^N.

Now assume that player 2's best-reply function is monotonic. Observe that our assumptions imply that $u_1(x^*) \geq u_1(x^N)$. Furthermore, we have

$$u_1(l_1, BR_2(l_1)) \geq u_1(x^N)\,.$$

Since x_1^* and l_1 must lie on two different sides of x_1^N, there exist two action profiles y and y' preferred by player 1 to x^N, but for which we have that $y_2 > x_2^N$ and $y_2' < x_2^N$ (because of the monotonicity of BR_2). This stands in contradiction with the initial assumption that player 1 has convex preferences and the single-valuedness of player 1's best-reply function. ∎

References

Admati, A. R., and M. Perry (1991): "Joint Projects without Commitment," *Review of Economic Studies*, 58, 259–276.

Bade, S., G. Haeringer, and L. Renou (2005): "More Strategies, more Nash equilibria," Forthcoming *Journal of Economic Theory*.

Bénabou, R., and J. Tirole (2004): "Willpower and Personal Rules," *Journal of Political Economy*, 112, 848–886.

Bernheim, B. D., and M. D. Whinston (1998): "Incomplete Contracts and Strategic Ambiguity," *American Economic Review*, 88, 902–932.

Clarke, F. H. (1989): *Methods of Dynamic and Nonsmooth Optimization*. CBMS-NSF Regional Conference Series in Applied Mathematics; 57, Society for Industrial and Applied Mathematics.

Conconi, P., and C. Perroni (2005): "Self-Enforcing International Agreements and Domestic Policy Credibility," mimeo, ECARES.

Dubey, P. (1986): "Inefficiency of Nash Equilibria," *Mathematics of Operations Research*, 11(1), 1–8.

Fudenberg, D., and J. Tirole (1991): *Game Theory*. MIT Press, Canbridge.

Gomez, A., and P. Jehiel (2005): "Dynamic Processes of Social and Economic Interactions," *Journal of Political Economy*, 113(3), 626–667.

Hamilton, J. H., and S. M. Slutsky (1990): "Endogeneous Timing in Duopoly Games: Stackelberg or Cournot Equilibria," *Games and Economic Behavior*, 2(1), 29–46.

Harris, C., P. J. Reny, and A. Robson (1995): "The Existence of Subgame-Perfect Equilibrium in Continuous Games with Almost Perfect Information: A Case for Extensive-form Correlation," *Econometrica*, 63, 507–544.

Harsanyi, J. (1973): "Oddness of the Number of Equilibrium Points: a New Proof," *International Journal of Game Theory*, 2, 235–250.

Hart, O., and J. Moore (2004): "Agreeing Now to Agree Later: Contracts that Rule Out but do not Rule In," Mimeo, London School of Economics.

Jackson, M., and S. Wilkie (2005): "Endogenous Games and Mechanisms: Side Payments among Players," *Review of Economic Studies*, 72, 543–566.

Lockwood, B., and J. P. Thomas (2002): "Gradualism and Irreversibility," *Review of Economic Studies*, 69, 339–356.

Maggi, G., and M. Morelli (2004): "Self-Enforcing Voting in International Organizations," Mimeo.

Moulin, H. (1984): "Dominance Solvability and Cournot Stability," *Mathematical Social Sciences*, 7(1), 83–102.

Osborne, M. J., and A. Rubinstein (1994): *A Course in Game Theory*. MIT Press.

Ray, D., and R. Vohra (1997): "Equilibrium Binding Agreements," *Journal of Economic Theory*, 73, 30–78.

Rockafellar, R. (1981): *The Theory of Subgradients and its Applications to Problems of Optimization. Convex and Nonconvex Funtions*. Research and Education in Mathematics, Heldermann Verlag Berlin.

Romano, R., and H. Yildirim (2005): "On the Endogeneity of Cournot-Nash and Stackelberg Equilibria: Games of Accumulation," *Journal of Economic Theory*, 120(1), 73–107.

Schelling, T. C. (1960): *The Strategy of Conflict*. Harvard University Press, Cambridge, MA.

van Damme, E., and S. Hurkens (1999): "Endogenous Stackelberg Leadership," *Games and Economic Behavior*, 28(1), 105–129.

Zeidler, E. (1986): *Nonliner Functional Analysis and its Applications I: Fixed-Point Theorems*. Springer-Verlag, New-York.