HWWA HAMBURG

DISCUSSION PAPER

Organizational Bargaining

Guillermo Owen

HWWA DISCUSSION PAPER

142

Hamburgisches Welt-Wirtschafts-Archiv (HWWA)
Hamburg Institute of International Economics
2000
ISSN 1616-4814

The HWWA is a member of:

- Wissenschaftsgemeinschaft Gottfried Wilhelm Leibniz (WGL)
- Arbeitsgemeinschaft deutscher wirtschaftswissenschaftlicher Forschungsinstitute (ARGE)
- Association d’Instituts Européens de Conjoncture Economique (AIECE)

Organizational Bargaining

Guillermo Owen

An earlier version of this paper was presented at the International Financial Markets (IFM) Seminar of the HWWA on July 2, 2001.

HWWA DISCUSSION PAPER

Edited by the Department
WORLD ECONOMY

Hamburgisches Welt-Wirtschafts-Archiv (HWWA)
Hamburg Institute of International Economics
Öffentlichkeitsarbeit
Neuer Jungfernstieg 21 - 20347 Hamburg
Telefon: 040/428 34 355
Telefax: 040/428 34 451
e-mail: hwwa@hwwa.de
Internet: http://www.hwwa.de/

Guillermo Owen
Department of Mathematics
Naval Postgraduate School
Monterey, California 93943
gowen@nps.navy.mil

Contents

Abstract

We consider a two-party bargaining model in which the two parties are organizations rather than individuals. Members of an organization have similar utilities for the agreements reached, but different utilities for conflict. We treat this as an *n*-person game in which the players are the members of the two organizations.

We look for the core of this game, and find that the core is always non-empty. In some cases, however (namely, when members of the two organizations are very hawk-like, or when the decision-making mechanisms are very rigorous), we find that the only core outcome ot the game is continued conflict.

Some discussion is included as to how the group leaders may be able to facilitate an agreement, and, conversely, as to how the most hawkish members of the organizations may be able to cooperate so as to bring about the unraveling of an agreement.

Organizational Bargaining

Guillermo Owen

1. The 2-Person Bargaining Model

A mathematical model for 2-person bargaining was first given by John Nash (1950) who considers the problem of two individuals, Adam and Barbara, who wish to divide a pie.

Nash posits a set K of possible agreements (the *feasible* set) as a compact, convex subset of the Euclidean plane. The coordinates of points represent utilities for the two players: essentially, a point $(u, v) \in K$ if there is a feasible agreement (division of the pie) between the two players, giving them utilities u and v respectively. Any point of K is available to them, if they can agree on it. The compactness property assures us that we can talk about maximal utilities (in particular, there will not be arbitrarily high utilities in the feasible set). Convexity is essentially guaranteed by the fact that the utilities in question are von Neumann-Morgenstern (1947) utilities: the utility of a lottery is equal to the expected utility of the lottery's outcome. Thus if (u, v) and (u', v') are two feasible points, any point on the line segment between them can be obtained as a lottery with these two points as outcomes.

Apart from this, Nash posits a *conflict point* (u_0 , v_0): if Adam and Barbara fail to reach an agreement, they will (with certainty) receive utilities u_0 and v_0 respectively. Nash assumes that $(u_0 , v_0) \in K$, as the two can always "agree to disagree."

Now, Nash defines a *bargaining game* as a triple, $[K, (u_0 , v_0)]$ satisfying the above conditions. Letting Ω be the space of all such triples, the *bargaining problem* consists in finding a function, defined on the space Ω, assigning to a given game $[K, (u_0 , v_0)]$ a point (u^*, v^*) which represents, in some sense, a fair outcome of the bargaining process (division of the pie).

Nash's approach is axiomatic: he gives a system of six axioms representing desirable (or necessary) properties of such a bargaining outcome.

The first axiom is *feasibility*: (u^*, v^*) must be a point of K. The second is *individual rationality*: $u^* \geq u_0$ and $v^* \geq v_0$.

Nash's third axiom is *Pareto-optimality,* defined as follow. We say that a point (u, v)$\in$ K is *Pareto-optimal* if there <u>does not exist another</u> (u', v') $\in$ K such that $u' \geq u$ and $v' \geq v$. Thus Pareto-optimal points are efficient points: neither player can increase his utility without decreasing that of the other. (Geometrically, the Pareto-optimal points of K form the "north-east" boundary of the set.) The third axiom simply says, then, that (u^*, v^*) must be a Pareto-optimal point of K. (See Figure 1)

Figure 1: A typical bargaining game

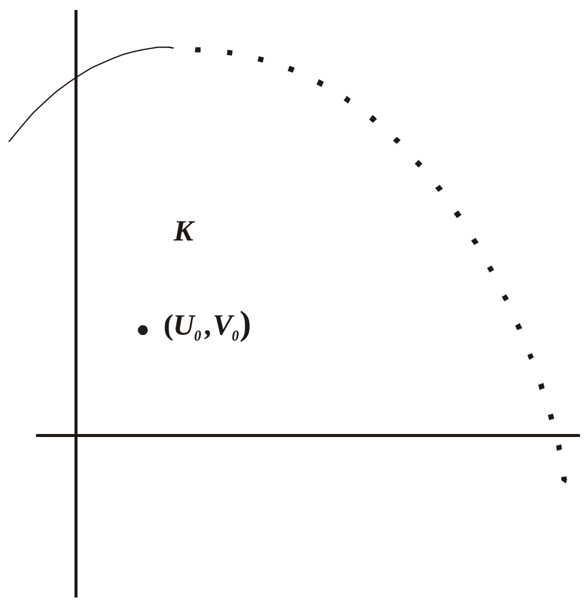

The Pareto-optimal curve is shown by · · · ·

These first three axioms seem eminently reasonable, and in general there is no quarrel with any of them. Nash's next three axioms are less obvious, and considerable discussion has arisen as to their reasonability; alternatives have been suggested by other researchers. We will not discuss them here, however, as they are not germane to the subject of this article.

2. Bargaining between Organizations

We consider, now, a situation in which the bargaining is no longer between the two individuals Adam and Barbara. Rather, we think of Adam and Barbara as agents for two organizations, A and B respectively. Thus, the situation becomes an *n*-person game, in which the players are the members of the two organizations. The main point to be made is that the several members of an organization have interests which are similar up to a point, but not coincident: the organizations are *heterogeneous.*

Once again, we assume that the feasible agreements can be represented by a compact convex set K in the Euclidean plane. We will however normalize this set by assuming that the Pareto-optimal frontier is a curve (possibly a straight line) passing through the points (0, 1) and (1, 0).

In the simplest of cases, the set K is given by the three inequalities, $u \geq 0$, $v \geq 0$ and $u + v \leq 1$. More generally, however, K will be given by the inequalities $u \geq 0$, $v \geq 0$, and $v \leq \varphi(u)$, where φ is a monotone decreasing, continuous, concave function, satisfying $\varphi(0) = 1$ and $\varphi(1) = 0$.

To model the similarities among members of an organization, we assume that their preferences *among agreements* are similar. Specifically, we assume that, if agreement (u, v) gives a bigger share of the pie to A than does (u', v'), then all members of A prefer (u, v). This is not supposed to mean that all members of A receive the same utility from an agreement, but only that their shares will (more or less) increase in some proportional way. (Note that in this paper we essentially deal with *ordina*l rather than *cardinal* utility; i.e., we talk about preferences but not about degrees of preference.)

The heterogeneity of the organizations will be represented by assuming that different members have different levels of tolerance (even desire) for conflict. One justification for this is that, in fact, it is not clear what exactly might happen if no agreement is reached. Thus, individual members of an organization might have different ideas as to how to value such an outcome. One member may feel that his organization is so strong that a continued conflict will eventually bring the other to its knees and thus lead to a very good outcome. Another justification is that the several members of an organization may have very different feelings about the rival organization. Thus an individual might hate the members of the opposite organization so much that he will only accept an agreement that gives a large amount to his own. In both of these cases, such individuals would generally be considered *hawks*.

On the other hand, there may be individuals who, either because they fear their organization is weak, or because they have a general distaste for fighting (or perhaps because they actually like members of the rival group), would accept almost any agreement. Such individuals would be considered *doves*.

To model this, we assume that each member, a, of A, can be assigned a number, $x(a)$, in the interval $0 \leq x \leq 1$. We say that individual a is located at position $x(a)$. Those individuals located near 0 will be called *doves*; those near 1 are *hawks*. Similarly, each individual member b of B is assigned a number $y(b)$, also in the interval $0 \leq y \leq 1$.

To model the actual degree of heterogeneity within organization A, a distribution function, F_A, is posited, satisfying

$$F_A(0) = 0,$$

$$F_A(1) = 1,$$

$$F_A(s) \leq F_A(s') \text{ whenever } s \leq s'$$

Here, we interpret $F_A(s)$ to be the fraction of organization A located at, or to the left, of s, i.e., those individuals a for whom $x(a) \leq s$. (Note that, when we speak of left – or right – in this context, we refer only to an individual´s location on the dove-hawk scale, and not to the usual political or economic distinction between right and left.)

A similar function F_B is assumed for organization B. See Figures 2 and 3 for possible examples of these functions. Note that the curves drawn do not actually represent the functions F_A or F_B; rather, they are the densities f_A and f_B; F_A is the integral of the given curve (the area under the curve f_B), and similarly for F_B.

Figure 2: A “dovish” density

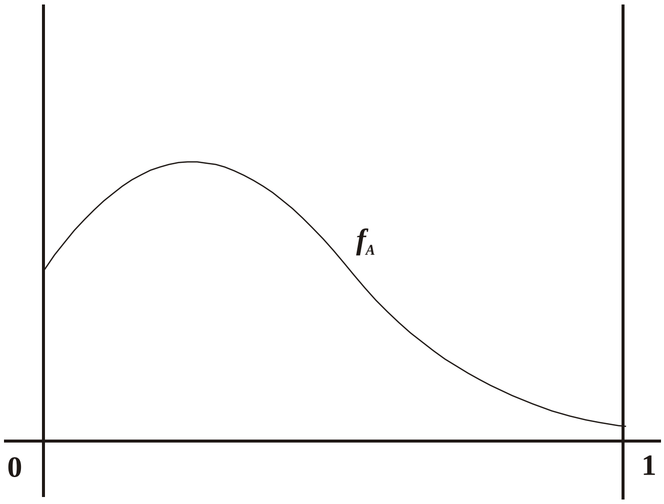

Figure 3: A “hawkish” density

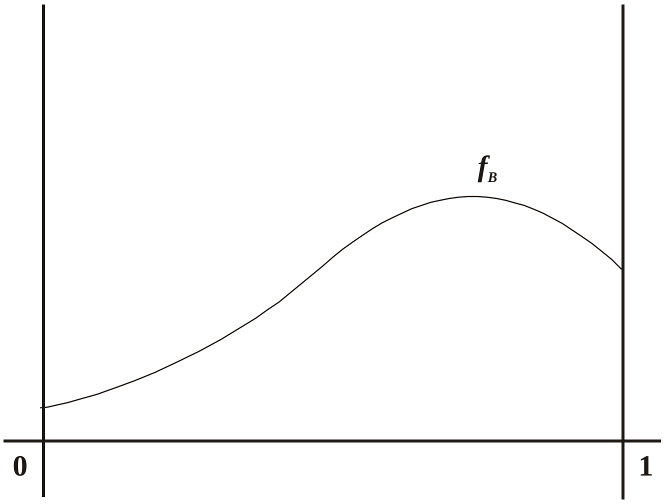

The *n*-Person Game

We will now develop this model so as to obtain an *n*-person cooperative game. Note that a game is defined by specifying (a) who the players are; (b) what possible outcomes might occur; (c) the preferences (utilities) which players have for these outcomes, and (d) the outcomes which the several coalitions (sets of players) can enforce.

In general, we will have here a game which can have two sorts of outcome: agreements, and continued conflict. As before, an agreement can be any point (u, v) in the set K. There is, however, only one continued conflict, which we will denote by C.

Now, the similarity among members of an organization lies in that they appreciate agreements similarly. Their dissimilarity lies in their taste for conflict. Thus, an individual member *a* of A, located at position $x(a)$, will have utility

$$U_a(u, v) = u$$

for an agreement, and

$$U_a(\mathrm{C}) = x(a)$$

for continued conflict. Similarly, an individual member *b* of B will have utility

$$U_b(u, v) = v$$

for an agreement, and

$$U_b(\mathrm{C}) = y(b)$$

for continued conflict.

We must next define the decision mechanism within each organization. *In the simplest case*, we shall assume the existence of a real number α, larger than ½ but smaller than 1, such that decisions in A must be approved by a fraction at least equal to α of its membership. Similarly, for B, we assume a number β, larger than ½ but smaller than 1, such that decisions in organization B must be approved by at least a fraction β of its members.[1]

The set $N = A \cup B$ is the set of all players. In other words, the players of this game are all the members of the two organizations. A *coalition* is any non-empty subset $S \subset N$. We distinguish three types of coalition. We say that a coalition S is *effective for* a given outcome if it can enforce that outcome, even against the concerted opposition of the complementary coalition N-S.

A *winning* coalition is any S containing *at least* a fraction α of A *and at least* a fraction β of B. Winning coalitions are effective for any outcome of the game.

A *blocking* coalition is any coalition which is not winning, but contains *either more than* a fraction 1-α of A, *or more than* a fraction 1-β of B. Blocking coalitions are effective for the outcome C (continued conflict) but not for agreements. (An alternative definition is to say that a blocking coalition is any coalition S which is not winning, but whose complement N-S is not winning.)

A *losing* coalition is any coalition that contains *not more than* 1-α of A, *and not more than* 1-β of B. Losing coalitions are not effective for any outcome. (An alternative definition is that a losing coalition is the complement of a winning coalition.)[2]

---

[1] Of course, the idea that decisions in either organization are made in such a democratic manner is probably too much of a simplification. More generally, we might assume decisions made in a more complicated manner. One possibility is that certain members of the organization might be considered more important than others. In such a case different weights might be assigned to them (a weight, e.g., of 100 for a head of government, 30 for the defense minister, 20 for other cabinet ministers or military leaders, 10 for parliamentarians, etc.), and then decisions would be made not by a fraction α of the total population, but rather by a fraction α of the total weights assigned. A more complicated rule would be to require, e.g., at least the president, at least 30% of the cabinet ministers, at least 50% of the army commanders, and at least 20% of the population. Clearly the possibilities are endless. This would not seriously change our method of approach. We would call such coalitions (subsets of A), *controlling coalitions for* A. In a similar way, controlling coalitions for B would be defined.

[2] Continuing with the ideas of footnote 1 above, a winning coalition would then be the union of two controlling coalitions, one for A and one for B. A losing coalition would be the complement of a winning coalition. A blocking coalition would be one which is neither winning nor losing. Again, winning coalitions are effective for any outcome, while blocking coalitions are effective only for the outcome C.

3. Characterization of the Core

The above determines an *n*-person game. We now analyze it by using the standard game-theoretic concept of the *core*. Originally defined in (Gillies, 1953), the core is one of the two or three most commonly used concepts for cooperative games.

In an *n*-person game, an outcome *X* is said to dominate another outcome, *Y*, if there exists a coalition *S*, all of whose members prefer *X* to *Y* (i.e., they all derive greater utility from *X*) and such that *S* is effective for *X*. The core is the set of all undominated outcomes.

Essentially, the core of a game represents its most stable outcomes. In many games, the question is whether there are any such outcomes; i.e., whether the core is non-empty. In fact, we shall see that, for the games here discussed, the core is always non-empty. However, this may not be as satisfactory a result as it seems at first view.

Figure 4

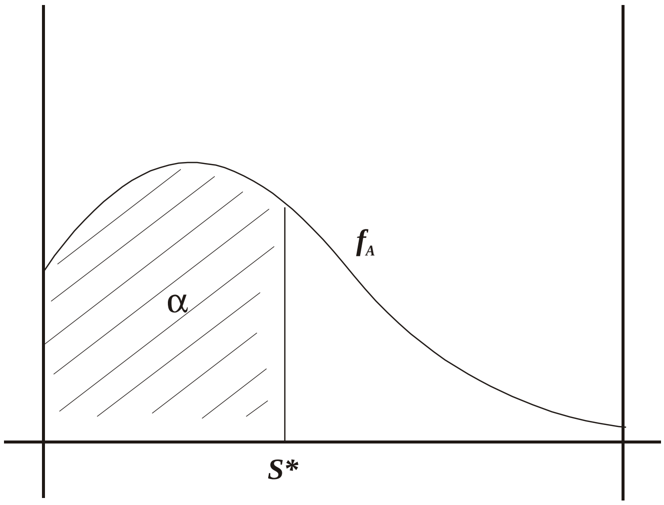

The core of the game will depend on the two numbers, s^* and t^*, defined below. The first is

$$s^* = F_A^{-1}(\alpha)$$

and is best defined as the smallest value of s such that $F_A(s) \geq \alpha$. Similarly,

$$t^* = F_B^{-1}(\beta)$$

and is the smallest value of t such that $F_B(t) \geq \beta$. (See Figures 4 and 5 for this.)[3]

Figure 5

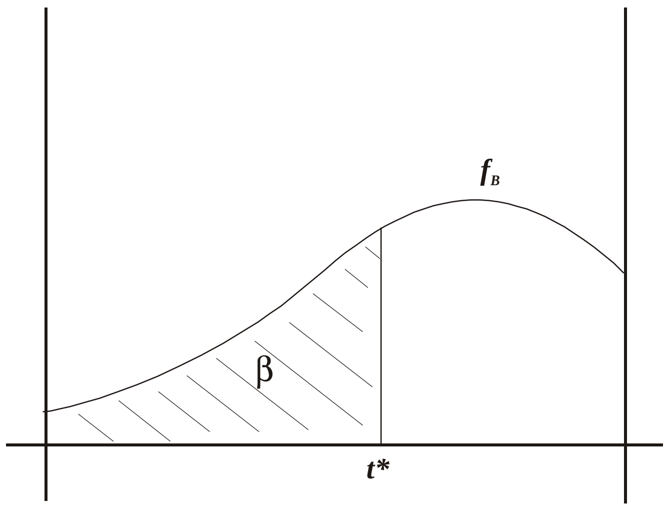

With these two defined, we are now able to describe the core of the game. It turns out that this depends on the position of the point (s^*, t^*) vis-à-vis the set K. Then, specifically,

[3] With the ideas of footnotes 1 and 2 in mind, we would define s^* as the smallest value of s, such that some controlling coalition (for A) is located at, or to the left of, s. The number t^* is defined similarly for organization B. With these definitions in mind, Theorem 1 continues to be true.

Theorem.

1. If (s^*, t^*) lies in the interior of the set K, then the core of the game consists of all Pareto-optimal points (u, v) such that $u \geq s^*$ and $v \geq t^*$. (Figure 6)

2. If (s^*, t^*) lies outside the set K, then the core of the game consists of the outcome C (continued conflict) only. (Figure 7)

3. If (s^*, t^*) is a Pareto-optimal point of K, then the core of the game consists of the agreement (s^*, t^*) and the outcome C. (Figure 8)

The proof of the theorem is given in the Appendix.

Figure 6: Case 1

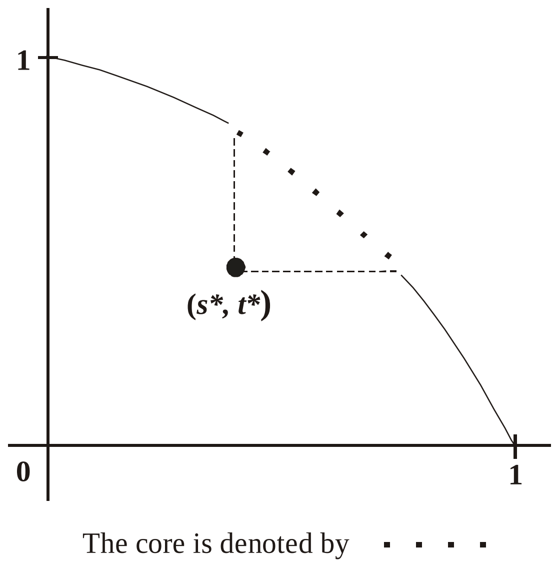

As may be seen, the non-emptiness of the core is not always very satisfying: in case 2 above, the core is (as always for these games) non-empty, yet we seem to be destined to continued conflict.

Figure 7: Case 2

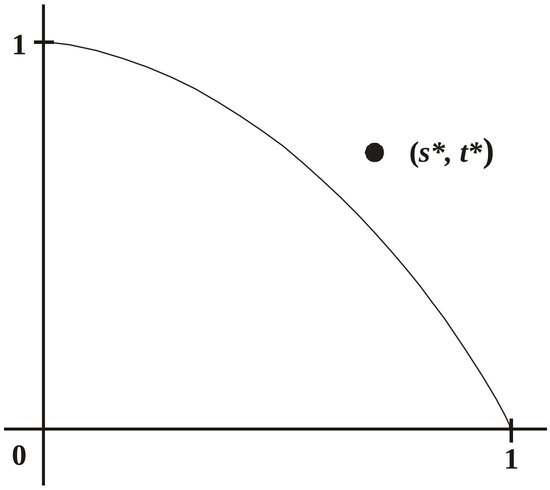

The core consists of the single outcome C.

Figure 8: Case 3

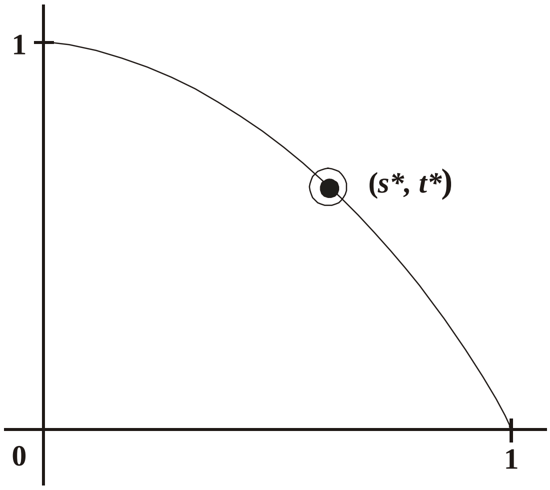

The core consists of the agreement
(s^*, t^*) and the conflict outcome *C*.

4. Discussion

To see exactly what the results of the theorem mean, note that (for a given K) case 1 will hold if s^* and t^* are small, while case 2 holds if these numbers are large (close to 1). Now, a large value for s^* means either that the population of organization A is very hawkish, or that a large super-majority is needed to obtain approval of an agreement. (Similar statements obviously hold for t^* and organization B.) Thus, if group leader Adam is having difficulty getting his people to agree, he will do best if he can move people leftward – i.e., cause them to appreciate peace more -- or else try to decrease α (the critical block size necessary for agreement).

In general, changing α is probably going to be very difficult, as it would mean changing the internal structure of the group. On the other hand, it may be possible to move people leftward into the peace camp. This can be done, possibly, by some sort of media propaganda campaign. In this, group leaders might explain the benefits to be expected from peace and warn of the dangers of continued conflict. Alternatively, it may be possible to depict members of the rival group as "reasonable" people. In this sense, a good media campaign might help to bring about an agreement, as the number s^* (or t^*) decreases.

An alternative possibility for bringing about an agreement would be for an interested third party to subsidize the agreement by "sweetening the pot," i. e., by making some sort of side payment to the members of the organization. Effectively, this can be represented as an increase in the set K. (See Figure 9)

Figure 9

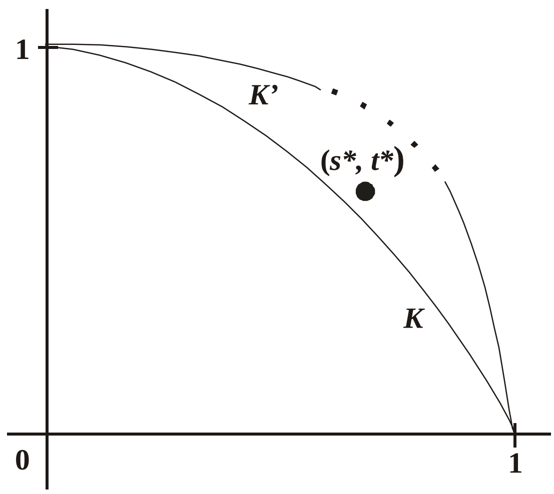

A side payment causes an increase in the set *K* (case 2) to the set *K'* (case 1) with core · · · · · · · ·

Conversely, consider the situation faced by someone who is unhappy with an agreement. For example, player *a*, an A member, will be unsatisfied with agreement (u, v) if $x(a) > u$; i.e., if he prefers continued conflict to the agreement. If so, he might try to undo the agreement by decreasing its utility, either to his own side A, or to side B. If he can get more than 1-α of side A, or more than 1-β of side B, to find the agreement unsatisfactory, he can then form a blocking coalition with these individuals.

Now, with members of his own organization, *a* can try to bring them to his point of view by political activity. With members of B, instead, he can try to decrease their utility (for an agreement) by taking action against them.

As an example, in the case of an armed conflict, suppose that – after an agreement has been reached – individual *a* – or, more likely, a group of like-minded individuals, all of them very hawkish members of A – decides to attack members of organization B by a car bomb, a suicide bomber, etc. In the case of labor strife, it may be that the more hawkish labor union activists decide to manhandle some of the replacement workers (strike-breakers) hired by management. In other conflicts, it may be nuisance lawsuits.

The effect of this action by (the hawks in) A will be to make members of B rethink the value of the agreement to them. In fact, the whole point of the agreement was that the organization would accept less than the best possible outcome so as to avoid further strife. If, however, it turns out that the agreement has not stopped the strife, then it is worth less than had been thought, and so a few of the marginal members of B may change their minds about acceptance. (Effectively, we can think of this as an increase in t^*.) Thus the number of members of B who support the agreement will decrease.

Nor is this all. After some violence by the recalcitrant members of A, some members of B may decide to follow suit, targeting A´s members and installations or otherwise creating problems for A. This will drive more members of A into the band of those opposed to the agreement. (Now it is s^* that increases.) Note that in this type of activity, the more hawkish members of the two organizations, though in theory bitter enemies, are in fact objective allies.

Figure 10

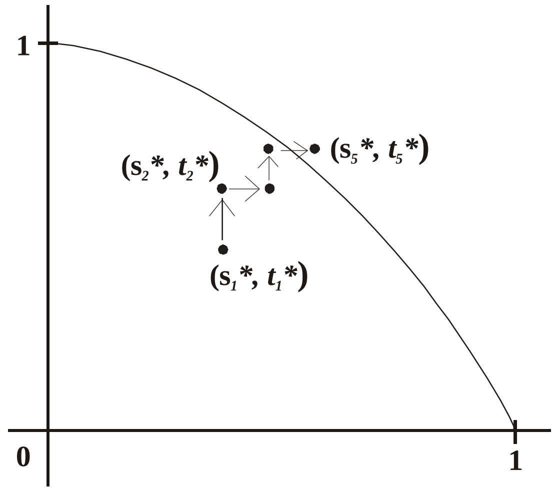

The agreement unravels (in four steps).

In practice, this type of activity can continue, with each side targeting the other, until the entire agreement unravels. (See Figure 10 for a graphic depiction of this.) To avoid this, some action must be taken. In general, Adam and Barbara must try to use their police powers (or whatever authority they have over their organizations) to rein in the wilder members. If they in fact do this quickly and efficiently, they can reassert their authority and save the agreement. If they do not take action quickly, Adam and Barbara will probably find themselves unable to maintain the treaty, and in fact may themselves be turned out of power (losing a vote of confidence, or perhaps worse if they represent violent organizations) if they insist on the agreement after it has fallen out of favor.

Appendix. Proof of theorem.

In proving this theorem, it is necessary to prove, on the one hand, that the given outcomes are undominated; on the other hand, that all other outcomes are dominated.

We prove first that all non-Pareto-optimal agreements are dominated by other agreements. In fact, suppose the agreement (u, v) is not Pareto-optimal. Then there exists some other agreement, $(u', v') \geq (u, v)$. Now, if in fact $u' > u$ and $v' > v$, then all players prefer (u', v'), and since N is always a winning coalition, (u', v') dominates (u, v).

There remains the possibility that either $u' > u$ and $v' = v$, or else $u' = u$ and $v' > v$. In either of these two cases, then by the fact that the function φ is continuous and strictly monotone decreasing, it will follow that there exists another agreement, (u'', v'') such that $u'' > u$ and $v'' > v$. Once again, we see that (u'', v'') dominates (u, v). We see thus that non-Pareto-optimal agreements cannot belong to the core.

Now, if (u, v) is Pareto-optimal, then for any other agreement (u', v'), it may be that $u' > u$, or it may be that $v' > v$, but not both. In the first case, (u', v') is preferred by all members of A but no members of B. Since A is not an effective coalition for an agreement, (u', v') does not dominate (u, v). In the second case, (u', v') is preferred by all of B but by none of A, and once again, since B is not effective for an agreement, (u', v') does not dominate (u, v).

Thus, a Pareto-optimal agreement cannot be dominated by any other agreement. Not all such outcomes, however, are in the core. The point is that they may still be dominated by the outcome C.

In fact, suppose (u, v) is an agreement, with $u < s^*$. We see (by definition of s^*) that $F_A(u) < \alpha$. In other words, more than $1-\alpha$ of organization A is located to the right of u, i.e., with $x(a) > u$. If we let S be this subset of group A, we see that all members of S prefer C to (u, v). Since S is a blocking coalition, this means C dominates (u, v). In a similar way, if $v < t^*$, C dominates (u, v) through some blocking $S \subset B$.

Suppose, however, that $u \geq s^*$ and $v \geq t^*$. Then $F_A(u) \geq \alpha$ and $F_B(v) \geq \beta$, and it follows that at most $1-\alpha$ of group A, and at most $1-\beta$ of group B, prefer C. Since this is not enough for a blocking coalition (and certainly not for a winning coalition), we see that C does not dominate (u, v). Thus, any Pareto-optimal (u, v) with $(u, v) \geq (s^*, t^*)$ is in the core of the game. Note that there are such agreements if, and only if, $(s^*, t^*) \in$ K.

It remains to determine whether C is in the core. Suppose first that (s^*, t^*) is an interior point of K. In that case, there will be at least one agreement (u, v) with $u > s^*$ and $v > t^*$. Now the agreement will be preferred to C by all $a \in$ A with $x(a) \leq s^*$ and by all $b \in$ B with $y(b) \leq t^*$. Thus the agreement is preferred by at least α of A and at least β of B. Since this is a winning coalition, the agreement dominates C.

Suppose, however, that (s^*, t^*) is a Pareto-optimal point of K. Then there is no feasible agreement with both $u > s^*$ and $v > t^*$. It will follow that an agreement can be preferred to C by a fraction α of group A or by a fraction β of group B, but not both. Thus, no winning coalition can be found to dominate C.

The same is true, *a fortiori*, when (s^*, t^*) lies outside the set K. C is then also undominated.[4]

[4] The proof of the theorem is slightly modified if we apply the ideas of footnotes 1, 2, and 3. The notation is slightly different, but the spirit of the proof is unchanged. We leave it to the more adventurous reader to modify the proof accordingly.

References

Gillies, D.B. (1953)
Some Theorems on *n*-Person Games. Ph. D.Thesis, Department of Mathematics, Princeton University.

McCormick, G. and G. Owen (1999)
Bargaining between Heterogeneous Coalitions. In: Mathematical and Computer Modelling.

Nash, J. (1950)
The Bargaining Problem. In: Econometrica.

Von Neumann, J. and O. Morgenstern (1947)
The Theory of Games and Economic Behavior. Princeton University Press.