

Bargaining and Majority Rules: A Collective Search Perspective*

Olivier Compte[†] and Philippe Jehiel[‡]

revised, February 2010

Abstract

We study a collective search process in which tentative proposals arrive sequentially and members of a committee decide whether to accept the current proposal or continue searching. The acceptance decision is made according to a (qualified) majority rule. We study which members have more impact on the decision, as well as the degree of randomness of the decision, as members get patient. When proposals vary along a single dimension, the acceptance set is small, and at most two members determine the outcome whatever the majority rule. When proposals vary along many dimensions, the acceptance set is large except under unanimity, and all members affect the distribution of decisions whatever the majority rule. Various implications are drawn.

1. Introduction

Many collective decisions take the form of a search process in which a committee examines new proposals sequentially and search stops when the current proposal receives the

*This paper has circulated under the title “Bargaining over randomly generated offers”. The authors are grateful to seminar audiences at University of Essex, Barcelona (JOCS), Paris (Roy), UCL (London), Stockholm, Yale, Princeton, University of Pennsylvania, NYU, Northwestern University, University of Chicago, Columbia. We also thank the editor and the anonymous reviewers for helpful comments that helped reshape the paper.

[†]PSE, 48 Boulevard Jourdan, 75014 Paris. e-mail:compte@enpc.fr

[‡]PSE and UCL; e-mail: jehiel@enpc.fr

support of sufficiently many committee members. Examples include recruitment decisions in which candidates are examined one by one, and the committee has to decide by vote whether to recruit now or pass this opportunity and wait for another candidate. In business decisions, capacity for financing projects is scarce. When a proposal or idea for a new project arrives, it is typically examined by a committee, and corporate culture affects the degree of consensus required for the proposal to be approved. In housing decisions, potential houses are examined sequentially, and the various members of the household have generally to agree for the transaction to take place.

We are interested in understanding how the majority requirement at the acceptance stage affects the outcome of *collective search problems* such as those described above. In particular, we focus on the following two questions.

- (1) Does every member of the committee have an effect on the set of possible agreements?
- (2) As members get patient, does the set of possible agreements get small?

The first question allows us to identify circumstances under which, despite having *formal* voting rights, some members would have *no real voting power* in the sense that small changes in the objectives or preferences of such members would not affect *at all* the set of possible agreements. The second question allows us to identify whether there are circumstances under which the decision process leads to *random outcomes* even when members are patient,¹ thus contrasting with individual search problems.²

As it turns out, it is not always the case that all members have real voting power, which is broadly consistent with a number of anecdotal evidence on electoral outcomes. This is also consistent with the prediction of the classic insight of the median voter theory (Downs (1957)): for all members but the median voter, small changes in preferences do not affect the identity of the median voter and the outcome is solely determined by the preference or bliss point of the median voter. The first contribution of our paper is to extend this classic insight to collective search environments and to examine how it

¹Impatience would inevitably translate into randomness in any search model whether collective or individual.

²In search problems in which a single individual decides, a patient individual would typically end up accepting only proposals that are close to the most preferred among the possible proposals, thereby making the search process almost deterministic in the limit.

generalizes to decision rules that differ from the simple majority rule (see the details of our results below).

Concerning the second question, we find that the set of possible agreements may remain large even as members get patient.³ Thus, the second contribution of our paper is to identify circumstances under which the *collective* character of the search process is itself a source of randomness.

The insights we derive will also imply that there are important differences between the predictions of the collective search model we examine, in which proposals put to a vote arrive randomly, and the collective bargaining model à la Baron and Ferejohn (1989) in which parties have a perfect control over the proposal put to a vote (see Section 5 of this paper for a comparison of these two models, and Section 6 for a model that unifies these two perspectives).

To be more specific, the answers to our two questions depend both on the *majority requirement* and the *dimensionality* of the proposal space (insofar as it affects the dimensionality of the utility space of members).

When proposals affect preferences along a *single* dimension (as in standard voting models), at most two members may have real voting power. We refer to these members as *key members*. The set of possible agreements (or agreement set) is determined by these key members in the sense that the same outcome would obtain if only the key members were present and unanimity among key members were required.⁴

Under simple majority with an odd number of members, only the median voter is a key member. Under more stringent majority requirement, there are two key members, and which two members are key depends on the majority rule, members' preferences and their relative impatience. Assuming that members differ only in their bliss points (and that members are risk averse), we find that under unanimity, the key members

³We also find that in such cases (and only in such cases), the distribution from which proposals are drawn may have a significant impact on the outcome.

⁴Members with no real voting power nevertheless have some effect on the outcome, insofar as they may affect the identity of the key members. For example, under simple majority, only the median voter is a key member, and the identity of the median voter clearly depends on the distribution of preferences.

are those with most extreme bliss points.⁵ Under weaker majority requirements, the agreement set is determined by two members (centered around the median voter) with less extreme bliss points.

As we move toward higher dimensional proposal spaces, we still get that not all members need have real voting power. But, when proposals affect members' preferences along as many dimensions as there are members, then all members affect the shape of the agreement set.⁶ In all cases we analyze, the agreement set when small appears to be close to the Nash bargaining solution restricted to the key members, i.e. the proposal that maximizes the product of the payoffs of the key members (when these are equally patient).

Concerning the size of the agreement, our findings are as follows. When proposals affect members' preferences along a single dimension, and agents have single-peaked preferences, the agreement set gets small as members get patient for any majority rule no less stringent than simple majority.

As we move toward higher dimensional proposal spaces, we still get that the agreement set gets small under unanimity. However, for any majority requirement other than unanimity, it remains of significant size when proposals affect members' preferences along sufficiently many dimensions. Thus, one should thus expect the decision to be more random under majority than under unanimity unless proposals affect members' preferences along a single dimension.

Extension to collective bargaining.

An important aspect of our model is that the proposals that are put to a vote are not controlled by the agents. In a number of international negotiations such as those prevailing at UNO, WTO or EU, the situation is neither one of pure collective search as parties may influence the proposal put to a vote nor one of pure collective bargaining as perfect control over the proposal put to a vote is difficult to achieve (as prior agreements or external news may create noisy inferences as to the implication of the proposal). This

⁵Members with bliss points less extreme than those of the key members do not affect the agreement set because they accept *all proposals* in the agreement set (and would even accept more proposals).

⁶In general, a member is key if there is proposal for which (i) he is pivotal and (ii) he is indifferent between adopting that proposal and waiting for another draw.

calls for extending our basic model to allow members to put some effort in influencing (possibly to their advantage) the draws over proposals, or in affecting the proposals' arrival rate.

In Section 6, we suggest how the insights derived in our basic collective search model in which no influence is possible are robust to such an extension. In particular, in all cases we have discussed so far, whenever the agreement set is small, the locus of the agreement set remains unchanged, so long as the control is not perfect (and the support of proposals does not change with the effort made by the members). This result shows the robustness of our insights about real voting power vs formal voting rights. We also suggest that whenever the agreement set is large in the collective search model, members have bigger incentives to generate proposals. Such a shift in incentives is desirable when, absent of effort, proposals arrive at a too slow pace, but undesirable otherwise because it may generate excessive and wasteful influence activities.

With international negotiations in mind, our model implies that whenever proposals vary along sufficiently many dimensions, a weaker majority requirement contributes in enhancing the role of the distribution of proposals, hence the scope for influence activities. Thus, in complex multi-issue international negotiations for which it is more likely that proposals affect members' preferences along many dimensions, our approach would predict a great deal of lobbying activities and an effect of all countries on the set of possible agreements. By contrast, in simpler single-issue international negotiations for which it is more likely that proposals affect members' preferences along fewer dimensions, our model suggests that formal voting rights need not translate into real voting power, and that the set of possible agreements tends to be determined by the countries with more extreme preferences when the majority requirement is stronger. Empirical work is needed to test these predictions.

Related literature.

Our paper can be viewed as contributing both to the literature on collective bargaining and to the literature on search. With respect to collective bargaining, the specificity of our model lies in relaxing agents' abilities to control the proposals put to a vote, as already highlighted. With respect to search, our main innovation is that the decision to

stop searching has to be approved by sufficiently many people as opposed to just one agent (see McCall 1970 for the pioneering paper or Rogerson et al. 2005 for a recent survey of the search literature). To the best of our knowledge, this paper is the first with Albrecht et al. (2009) to examine collective search models under arbitrary majority rules.⁷ While our focus is on understanding the factors that drive the final outcome, Albrecht et al.'s main concern is on comparing how the collective search problem differs from a single agent search problem and in so doing Albrecht et al. restrict attention to the case of symmetric agents whereas no such restriction is being made in our analysis. The predecessor of this paper (Compte and Jehiel 2004) and Albrecht et al. intersects on the derivation of the insight that in symmetric setups ex ante welfare is optimized with more stringent majority requirements, as agents get more patient.⁸

2. The Model

We consider a committee consisting of n members, labeled $i = 1, \dots, n$.

Timing. At any date $t = 1, \dots$, if a decision has not been made yet, a new proposal is drawn and examined. We denote by x a proposal, by X the set of proposals. Proposals at the various dates $t = 1, \dots$ are drawn independently from the same distribution with density $f(\cdot) \in \Delta(X)$.

Upon arrival of a new proposal x , each member decides whether to accept that proposal. We consider various majority rules. Under the k -majority rule, the game stops whenever at least k out of the n members vote in favor of the proposal.

Proposals. The set of proposals X is assumed to be a compact convex subset of

⁷For the case of unanimity rule and two agents, Wilson (2001) is a precursor of this paper. He obtains for this case the link to the Nash bargaining solution as parties get very patient. In the discussion section of his paper, Wilson also suggests how his insight could be extended to the more than two player case under additional assumptions on the distribution of offers. The insight we obtain in the unanimity rule case for the case of rich proposal space generalizes Wilson's insight for the case of more than two players, in particular revealing that there is no need for further assumption on the distribution of offers.

⁸In a recent paper, Strulovici (2007) considers the issue of experimentation by voting. He extends the literature on experimentation in the same way as we extend the search literature by considering a multi-agent setup with voting procedures rather than a single agent setup.

\mathfrak{R}^m with $1 \leq m \leq n$, and we shall refer to m as the dimension of the proposal space. Throughout the paper, we also assume that the density f is continuously differentiable and bounded away from 0 on X . We call \mathcal{F} the set of such densities.⁹

Preferences. We let $u_i(x)$ denote the utility that member i derives from decision x at the time it is implemented. We normalize to 0 the payoff that any member i obtains under perpetual disagreement. Throughout the paper, we will assume that all proposals are individually rational for all players, that is, $u_i(x) > 0$ for all $i \in \{1, \dots, n\}$. We shall also assume that each u_i is concave, continuously differentiable on X , and locally non-constant.¹⁰

Member i discounts future payoffs according to a discount factor δ_i , so that viewed from date 1, a decision x agreed upon at date t yields member i a discounted payoff equal to $(\delta_i)^{t-1} u_i(x)$.

Richness of proposal space. In general, the dimension m of the proposal space X need not coincide with the dimension of the set of utility vectors $u(X)$. This is, for example, the case when there are only few dimensions of x that members care about.¹¹ Throughout the paper, we take the convention that members care about all dimensions of the proposals so that the dimension of proposal space also reflects the dimension of $u(X)$ in the utility space. Technically, we shall assume that the rank of the gradient matrix of u has everywhere rank m .

Strategies and equilibrium. In principle, a strategy specifies an acceptance rule that may at each date be any function of the history of the game. We will however restrict our attention to *stationary* equilibria of this game, where the acceptance rules adopted

⁹Specifically, given the proposal space X , we shall restrict attention to densities in $\mathcal{F} = \{f \in C^1(X), f(\cdot) \geq b \text{ and } |\frac{\partial f}{\partial x_j}| < d \text{ on } X \text{ for all } j = 1, \dots, m\}$ for some $b > 0$ and $d > 0$.

¹⁰We say that u is locally non-constant if the following two conditions hold: (i) For any $\varepsilon > 0$, there exists $\varepsilon' > 0$ such that for any convex set $A \subseteq X$ of size larger than ε (see (2.5) for the definition of the size), the set $u(A) = \{u \in \mathfrak{R}^n, \exists x \in A, u_i = u_i(x) \text{ for all } i\}$ has size larger than ε' . (ii) For any $\varepsilon > 0$, there exists $\varepsilon' > 0$ such that for any convex set $A \subseteq X$ of measure larger than ε , the sets $u_i(A)$ have size larger than ε' for all i .

¹¹Technically, the set $u(X) = \{u = (u_1, \dots, u_n) \mid \exists x \in X, u_i(x) = u_i, \forall i\}$ may lie on a manifold of dimension lower than m .

by the various members do not depend on the calendar time t .¹²

Given any stationary acceptance rule σ_{-i} followed by members j , $j \neq i$, we define the expected payoff $\bar{v}_i(\sigma_{-i})$ that member i derives given σ_{-i} by following his (best) strategy. An optimal acceptance rule for member i is thus to accept the proposal x if and only if¹³

$$u_i(x) \geq \delta_i \bar{v}_i(\sigma_{-i}),$$

which is stationary as well (this defines the best-response of member i to σ_{-i}).

Stationary equilibrium acceptance rules are thus characterized by a vector $v = (v_1, \dots, v_n)$ such that member i votes in favor of x if $u_i(x) \geq \delta_i v_i$ and votes against it otherwise. For any k -majority rule and value vector v , it will be convenient to refer to $A_{v,k}$ as the corresponding acceptance set, that is, the set of proposals that get support from at least k members when failing to agree today yields member i a continuation payoff of v_i (from the viewpoint of next period).¹⁴

$$A_{v,k} = \{x \in X, \exists K \subset \{1, \dots, n\}, |K| = k, u_i(x) \geq \delta_i v_i \text{ for all } i \in K\}. \quad (2.1)$$

Equilibrium consistency then requires that

$$v_i = \Pr(x \in A_{v,k}) E[u_i(x) \mid x \in A_{v,k}] + [1 - \Pr(x \in A_{v,k})] \delta_i v_i \quad (2.2)$$

or equivalently

$$\delta_i v_i = \lambda_i E[u_i(x) \mid x \in A_{v,k}] \quad (2.3)$$

where

$$\lambda_i = \frac{\Pr(x \in A_{v,k})}{1 - \delta_i + \delta_i \Pr(x \in A_{v,k})}. \quad (2.4)$$

¹²To avoid coordination problems that are common in voting (for example, all players always voting "no"), we will also restrict attention to equilibria that employ no weakly dominated strategies (in the stage game). These coordination problems could alternatively be avoided by assuming that votes are sequential.

¹³Strictly speaking, member i 's behavior is indeterminate when $u_i(x) = \delta_i \bar{v}_i(\sigma_{-i})$. Yet, cases of indifferences, i.e. such that $u_i(x) = \delta_i \bar{v}_i(\sigma_{-i})$, are irrelevant, as they have measure 0 (according to $f(\cdot)$).

¹⁴For any finite set B , $|B|$ denotes the cardinality of B .

A stationary equilibrium is characterized by a vector v that satisfies (2.1)-(2.2). We first observe that such an equilibrium always exists.

Proposition 1 (*Existence*) *Whatever the majority requirement, a stationary equilibrium exists.*

Proof: Define the function $v \rightarrow \phi(v)$, where $\phi_i(v)$ coincides with the RHS of Equation (2.2), and let $\bar{u} = \max_{i,x} u_i(x)$. The function ϕ is continuous from $[0, \bar{u}]^n$ to itself because $u(\cdot)$ is locally non-constant. Hence, ϕ has a fixed point, which is a stationary equilibrium. **Q. E. D.**

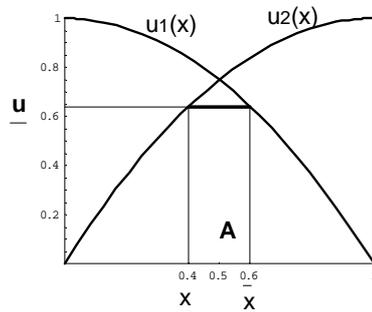
In the rest of the paper, we will be addressing two questions. First, we will be concerned with the size of the agreement set A as players get more and more patient where the size of A is defined as

$$size(A) = \sup_{(x,y) \in A^2} \|x - y\|, \quad (2.5)$$

and $\|x - y\|$ denotes the euclidean distance between x and y . Second, we will investigate whether and when the preferences of some players are (locally) irrelevant for the determination of the agreement set A .

Before we get into the heart of the results, we illustrate how to construct a stationary equilibrium in two simple two-member examples.

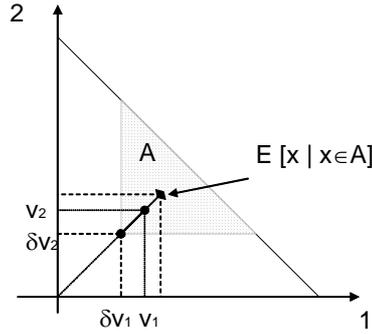
A one-dimensional example. We consider the following two-player example where (i) decisions require unanimity; (ii) proposals are drawn uniformly from the unit interval $X = [0, 1]$; (iii) member 1's preference is $u_1(x) = 1 - x^2$ and member 2's preference is $u_2(x) = 1 - (1 - x)^2$; (iv) members have the same discount factor δ .



As shown in the above Figure, the acceptance set takes the form of an interval centered around $\frac{1}{2}$, i.e. $A = (\frac{1}{2} - z, \frac{1}{2} + z)$ where z is such that member 1 (resp. 2) is indifferent between accepting $\bar{x} = \frac{1}{2} + z$ (resp. $\underline{x} = 1 - z$) now and postponing the decision. Given that $\Pr(A) = 2z$ and $E(u_1(x) \mid x \in A) = \frac{3}{4} - \frac{z^2}{3}$, the equilibrium requirement (2.2) boils down to finding z and an acceptance threshold \underline{u} such that

$$\underline{u} = u_1(\frac{1}{2} + z) = 1 - (\frac{1}{2} + z)^2 = \frac{2z\delta}{1 - \delta + 2z\delta}(\frac{3}{4} - \frac{z^2}{3}).$$

A rich type-space example. We consider the following example where (i) decisions require unanimity; (ii) proposals are drawn uniformly on the simplex $X = \{(x_1, x_2), 0 \leq x_i, x_1 + x_2 \leq 1\}$; (iii) preferences are linear, i.e. $u_i(x) = x_i$; (iv) members have the same discount factor δ .



As shown in the above Figure, the acceptance set takes the form of a small simplex $A = \{x = (x_1, x_2) \mid x_i \geq \delta v_i \text{ for } i = 1, 2\}$ where each v_i is a weighted average between $E[x_i \mid x \in A]$ and δv_i , with a weight on the former equal to $\Pr(A)$. Since $E[x_i - \delta v_i \mid x \in A] = \frac{1}{3}[1 - \delta v_1 - \delta v_2]$ and $\Pr A = (1 - \delta v_1 - \delta v_2)^2$, the equilibrium requirement (2.2) boils down to finding v and an acceptance threshold $\underline{u} = \delta v$ such that $\frac{1-\delta}{\delta}\underline{u} = \frac{1}{3}(1 - 2\underline{u})^3$.

3. When proposals vary along a single dimension.

In this Section, we assume that proposals vary along a single dimension $X = [0, 1]$, which is well suited to deal with collective decisions bearing on single issues, or when there is a single dimension that members care about. We start by assuming that members differ

only in their bliss point. That is, we assume that all members have the same discount factor δ and that the utility of member i with bliss point θ_i writes $u_i(x) = v(x - \theta_i)$ for some concave function v assumed to be smooth, single-peaked with a maximum at 0. Reordering members by increasing order of bliss points, and recalling that any proposal is preferable to the status quo, our assumptions can be summarized into:

Assumption 1: *Assume $X = [0, 1]$, $0 \leq \theta_1 < \dots < \theta_n \leq 1$, $u_i(x) = v(x - \theta_i)$, where v is smooth, single-peaked with a maximum at 0, concave and positive on $[-1, 1]$.*

We start with a result concerning the size of the agreement set. We then characterize who drives the outcome as a function of the majority rule. We also characterize the locus of the agreement set as players are very patient.

3.1. Size of the agreement set.

Our first result shows that whatever the majority requirements $k > n/2$, in the limit as members are very patient, the agreement set is a small interval.

Proposition 2: *Consider a majority requirement k . In any equilibrium, the agreement set is an interval. Besides, for any majority requirement ($k > n/2$), and for any $\varepsilon > 0$, there exists δ_0 such that for all $f \in F$ and $\delta \geq \delta_0$, any equilibrium agreement set has size below ε .*

Intuitively, the reason why the agreement set is small is the following. If the agreement set is large, then there are essentially no delay costs when δ is close to 1. This implies that every member should reject the proposal in A he likes least, which given the single-peakedness of preferences is either one or both extreme points of the agreement set. Thus, either one or both extreme points of the agreement set would be rejected by at least $n - k$ members (given that $k > \frac{n}{2}$), thereby leading to a contradiction to the very definition of the agreement set. The detailed proof appears in Appendix.

Proposition 2 has implications regarding the *sensitivity* of the equilibrium analysis with respect to the distribution of proposals $f(\cdot)$. Whenever the agreement set A is

small for some density f_0 , then $E[u_i(x) \mid x \in A]$ would vary little if the density were to switch to $f(\cdot)$. This in turn guarantees (in a way that we make precise in subsection 4.1) that the density $f(\cdot)$ has little effect on equilibrium outcomes. One implication of this observation will be that in a setting where players would have the option to affect the density $f(\cdot)$ through influence activities, there would be little benefit in doing so (see Section 6).

3.2. Who drives the outcome.

As shown in Proposition 2, the agreement set is an interval $A = [\underline{x}, \bar{x}]$. The next result shows that the two boundary points of the agreement set are determined by the preferences of (at most) two members, that we shall refer to as *key members*, and it also characterizes who these key members are. Note that the characterization holds whether members are patient or not.

Proposition 3: *Fix a majority rule $k > n/2$. Consider an equilibrium with agreement set $A = [\underline{x}, \bar{x}]$, and equilibrium values $v = (v_i)_{i \in I}$. Let $i_0 = n - k + 1$ and $i_1 = k$ and $M = \{i, i_0 < i < i_1\}$. We have:*

- (i) $u_{i_0}(\bar{x}) = \delta v_{i_0}$ and $u_{i_1}(\underline{x}) = \delta v_{i_1}$.
- (ii) For all members $i \in M$, $u_i(x) > \delta v_i$ for all $x \in A$.
- (iii) The game where unanimity is required among members i_0 and i_1 only also has A as an equilibrium agreement set.

The members i_0 and i_1 defined in Proposition 3 are the key members. Under the simple majority rule with an odd number of players, there is a single key member ($i_0 = i_1 = \frac{n+1}{2}$) who coincides with the median member. In all other cases, there are two distinct key members, whose preferences are more extreme as the majority requirement is increased. Under unanimity, the key members are those members with most extreme preferences.

Observations (i) and (ii) of Proposition 3 show that member i_0 (respectively i_1) is indifferent between accepting and rejecting proposal \bar{x} (respectively proposal \underline{x}), and

that all members that are intermediate between i_0 and i_1 accept any offer in A . Observation (iii) shows the precise sense in which members i_0 and i_1 are key. Members i_0 and i_1 are key in the sense small changes in the preferences of members other than i_0 and i_1 would not affect the outcome.^{15,16}

To get some intuition for Proposition 3, consider two members i_0 and $i_1 \geq i_0$, and consider the game where unanimity is required among these two members i_0 and i_1 only, along with an equilibrium agreement set A of that game. Our result is driven by the observation that members with intermediate bliss point (i.e. such that $i_0 < i < i_1$) have preferences on A that are flatter than those of members i_0 and i_1 . As a result, they are better off accepting any proposal in A rather than waiting for a better draw in A . The precise location of their bliss point θ_i within the interval $(\theta_{i_0}, \theta_{i_1})$ is irrelevant.

In contrast, any member with preferences that are more extreme than those of i_0 or i_1 cares about *which* element x in A is drawn, and would be willing to reject those draws that are furthest away from his bliss point. He may thus affect A , but only to the extent that he is pivotal, and this depends on the majority rule. In equilibrium, the key members i_0 and i_1 are precisely chosen so that any member more extreme than i_0 or i_1 is not pivotal for those draws he wants to reject: whether he rejects such draws or not, there is still a k -majority that favors them.

3.3. The locus of the agreement set.

We use previous results to derive the location of the agreement set considering the limit of very patient members. Define $x_{i,j}^*$ as the Nash bargaining solution between members i and j . That is,¹⁷

$$x_{i,j}^* = \arg \max_x v(x - \theta_i) \cdot v(x - \theta_j).$$

¹⁵Small changes in the preferences of non-key members would not alter the ordering of members' blisspoints, hence the identity of i_0 and i_1 , and by (iii) the agreement set is solely determined by the preferences of i_0 and i_1 .

¹⁶Note however that the presence of members other than i_0 and i_1 has some effect on the outcome to the extent that these members contribute to determining the identity of the key members.

¹⁷Observe that due to the concavity of $v(\cdot)$, the Pareto frontier in the space of members i and j 's preferences is convex as one varies x .

Note that if v is symmetric around 0, then $x_{i,j}^* = \frac{\theta_i + \theta_j}{2}$. Also note that $x_{i,i}^* = \theta_i$.

We show that, irrespective of the density $f \in F$, the solution is determined by the Nash bargaining solution among the two members $i_0 \equiv 1 + n - k$ and $i_1 \equiv k$ defined in Proposition 3.¹⁸

Proposition 4: *Let Assumption 1 hold, and let $k > n/2$ be the qualified majority rule. Then for any density $f \in \mathcal{F}$, when δ tends to 1, only proposals close to $x^* = x_{1+n-k,k}^*$ are accepted.*

Intuitively, at the limit where members are very patient, the agreement set shrinks to a small neighborhood of some $x^* \in (\theta_{i_0}, \theta_{i_1})$. We characterize x^* using the two indifference conditions that appear Proposition 3 (i),¹⁹ and we note that x^* does not depend on the distribution over proposals. Specifically, defining

$$\mu_i(x) = \frac{|u'_i(x)|}{u_i(x)} \quad (3.1)$$

as member i 's *intensity of preferences* at x , we show that x^* satisfies:

$$\mu_{i_0}(x^*) = \mu_{i_1}(x^*) \quad (3.2)$$

hence the connection with the First Order Conditions characterizing the Nash bargaining solution. Echoing observation (ii) in Proposition 3, note that for any member i with bliss point $\theta_i \in (\theta_{i_0}, \theta_{i_1})$, we have

$$\mu_i(x^*) < \mu_{i_0}(x^*)$$

which is another way of explaining why members with bliss points $\theta_i \in (\theta_{i_0}, \theta_{i_1})$ are willing to accept any proposal falling in A (compared to i_0 and i_1 , they have flatter preferences over A).

¹⁸We provide a *direct* proof of Proposition 4 in the Appendix. Yet, given Proposition 3, which says that the equilibrium outcome must be the same as the one obtained under unanimity when only members $i_0 \equiv 1 + n - k$ and $i_1 \equiv k$ are present, Proposition 4 can also be viewed as a corollary of our analysis of the rich proposal space case, assuming for that purpose that there are only two members, i_0 and i_1 (see Section 4). It can also be viewed as deriving from Wilson (2001)'s analysis.

¹⁹These two indifference conditions are $u_{i_0}(\bar{x}) = \lambda E[u_{i_0}(x) \mid x \in A]$ and $u_{i_1}(\underline{x}) = \lambda E[u_{i_1}(x) \mid x \in A]$.

3.4. Discussion

3.4.1. Comparing majority and unanimity rules.²⁰

The above propositions show how we should expect the set of accepted proposals to vary with the majority requirement. We first assume that $v(\cdot)$ is symmetric around 0 and consider the effect of the distribution over bliss points. We next consider the effect of asymmetric $v(\cdot)$, assuming that bliss points are evenly distributed.

When the function $v(\cdot)$ is symmetric the unanimity rule leads to outcomes close to $\frac{\theta_1 + \theta_n}{2}$ and the k -majority rule leads to outcomes close to $\frac{\theta_{n-k+1} + \theta_k}{2}$. If bliss points are evenly distributed between 0 and 1, no much difference should be expected, as one varies the majority requirement.²¹ However, if the distribution of bliss points is not even, say it is skewed toward 0, then while unanimity leads to outcomes that do not reflect this skewedness (if $\theta_1 = -1$ and $\theta_n = 1$, outcomes close to $\frac{1}{2}$ are obtained irrespective of the distributions of other bliss points), less stringent majority rules lead to outcomes more biased toward 0.

>From a more statistical perspective, assume the bliss points are randomly drawn. Specifically, assume that there are three groups of players, N_0 , N_+ and N_- , of random size n_0 , n_+ and n_- . Further assume that members in group N_0 , N_+ and N_- have bliss points drawn in a neighborhood of $\frac{1}{2}$, θ^+ and θ^- , respectively. Which majority requirement induces more volatile decisions? Who benefits from which majority requirement?

Under unanimity, the agreement will be in a neighborhood of $\frac{\theta^+ + \theta^-}{2}$. Members of N_+ will benefit from having a less homogenous group, because this tilts the outcome towards proposals that they prefer.

As the majority requirement is reduced, and so long as n_0 is large enough, the outcome will get close to one that players in N_0 favor. If however n_0 is small, then reducing the majority requirement may lead to a large volatility in the agreement, as it

²⁰We are grateful to the editor for suggesting the insights developed in this subsection.

²¹Observe that even if the distribution of bliss points is even, there would be a notable effect of the majority requirement if the function $v(\cdot)$ were assumed to be asymmetric around 0. For example, if $v'(z) > v'(-z)$ for all $z > 0$, then $x^{NB}(k, r) > \frac{\theta_k + \theta_r}{2}$. Accordingly, assuming bliss points are evenly distributed, the unanimity rule yields outcomes more biased toward 1 as compared with the simple majority rule.

may go from neighborhoods of θ^+ to neighborhoods of θ^- as a function of the realized difference $n_+ - n_-$.

We next observe that asymmetries in the function $v(\cdot)$ may be another source of discrepancy between majority and unanimity.. Consider for example an even distribution of bliss points on $[0, 1]$, and assume that for any $a > 0$, $v(a) < v(-a)$. Under the simple majority rule, asymmetries are *irrelevant*: the outcome coincides with the median bliss point, even if the loss $v(0) - v(a)$ from adopting a proposal $\theta_i + a$ lying above one's bliss point very much exceeds the loss $v(0) - v(-a)$ from adopting a proposal $\theta_i - a$ lying below one's bliss point. As the majority requirement increases however, the outcome will reflect that difference in losses: being determined by the Nash bargaining solution between two members, the outcome will be driven towards 1.

3.4.2. Other forms of heterogeneity

So far we have assumed that members differ only in their bliss point. In this subsection we consider how our results are altered when other forms of heterogeneity are considered.

To fix ideas, suppose first that members differ only in their degree of impatience but are otherwise identical. Then it is easy to see that there is a single key member, in the sense that the outcome is the same as if only this member were searching for a proposal: under unanimity, the key member is the most patient agent; under the k -majority rule, the key member is the agent i with the $(n - k + 1)$ highest discount factor δ_i . In other words, the identity of the key member depends on the relative impatience of the members together with the majority requirement.

We now generalize our previous results to cases where members are heterogeneous *both* in their bliss points and in their discount factor. Preferences are as assumed above (see Assumption 1) and we let δ_i be such that $1 - \delta_i = \frac{1-\delta}{\alpha_i}$ for some positive α_i . Much of our previous results did not rely on the discount factor being common. It is easy to check that the agreement set is an interval that vanishes when δ tends to 1 (Proposition 2) and that the interval must be determined by the preferences of two key members (Proposition 3 (iii)).

The only change is what determines *the identity of the key members*. Once the

key members are determined, the locus of the agreement set will correspond (as in Proposition 4) to a Nash bargaining solution between these key members, appropriately modified to take into account the possible asymmetry in the discount factors.

To characterize the key members, it is convenient to define, for each $x \in (\theta_1, \theta_n)$ and member i ,

$$\mu_i(x) \equiv \alpha_i \frac{|u'_i(x)|}{u_i(x)},$$

as a measure of member i 's *intensity of preferences at x* . This definition generalizes (3.1) and allows members to be ranked according to the intensity of their preferences, from the more moderate (low μ_i) to the more extreme (high μ_i), giving highest rank to the most extreme member.²²

Consider next the group of members for which $\theta_i > x$ (respectively $\theta_i \leq x$). Within that group, we define the member $i_1^{(m)}(x)$ (respectively $i_0^{(m)}(x)$) having the m^{th} highest rank, and denote by $\mu_+^{(m)}(x)$ (respectively $\mu_-^{(m)}(x)$) the intensity of his preference at x .²³ For each $m \leq n/2$, there is a unique solution to the equation:²⁴

$$\mu_+^{(m)}(x) = \mu_-^{(m)}(x).$$

We denote by $x^{(m)}$ that solution, and let $i_0^{(m)} \equiv i_0^{(m)}(x^{(m)})$ and $i_1^{(m)} \equiv i_1^{(m)}(x^{(m)})$.²⁵

By construction, the point $x^{(m)}$ corresponds to the asymmetric Nash bargaining solution between the two members with weights given by their relative patience. Formally,

²²Note that in the case of heterogeneous α , the ranking generally depends on x . Also note that two members may have the same intensity of preferences for some x . In that case, we take the convention of ranking higher the individual with highest index.

²³If there are fewer than m members such that $\theta_i > x$ (resp. $\theta_i < x$), the member $i_1^{(m)}(x)$ (resp. $i_0^{(m)}(x)$) is not defined, but we let $\mu_+^{(m)}(x) = 0$ (resp. $\mu_-^{(m)}(x) = 0$).

²⁴Because of concavity $\mu_-^{(m)}(x)$ is increasing in x (from 0), while $\mu_+^{(m)}(x)$ is decreasing in x (down to 0). Since these functions are continuous, and since for $m \leq n/2$, there exists x such that $\mu_-^{(m)}(x) > 0$ and $\mu_+^{(m)}(x) > 0$, there is a single solution $x^{(m)}$ and it satisfies $\mu_-^{(m)}(x^{(m)}) > 0$, so $i_0^{(m)}$ and $i_1^{(m)}$ are well defined at $x^{(m)}$.

²⁵By construction, the proposal $x^{(m)}$ and the two members $i_0^{(m)}$ and $i_1^{(m)}$ have the property that at $x^{(m)}$: (1) members $i_0^{(m)}$ and $i_1^{(m)}$ have the same intensity of preferences; (2) there are exactly $m - 1$ members with bliss point below $x^{(m)}$ who are more extreme than $i_0^{(m)}$ and $m - 1$ members with bliss point above $x^{(m)}$ who are more extreme than $i_1^{(m)}$; (3) the other members are at least as moderate as $i_0^{(m)}$ and $i_1^{(m)}$ at $x^{(m)}$.

for any two members i and j , define

$$x_{i,j}^{ANB} = \arg \max_x v(x - \theta_i)^{\alpha_i} v(x - \theta_j)^{\alpha_j}$$

The point $x^{(m)}$ coincides with $x_{i,j}^{ANB}$ with $i = i_0^{(m)}$ and $j = i_1^{(m)}$.

The following proposition shows that under the k -majority rule, the key members are $i_0^{(m)}$ and $i_1^{(m)}$ with $m = n + 1 - k$ and only proposals close to $x^{(m)}$ are accepted.

Proposition 5: *Let Assumption 1 hold. Consider the k -majority rule with $k > n/2$. When δ tends to 1, the agreement set is the same as the one when unanimity prevails and only members $i_0^{(1+n-k)}$ and $i_1^{(1+n-k)}$ are present. Moreover, only proposals close to $x^{(1+n-k)}$ get accepted.*

To sum up, our insight that at most two members determine the outcome when proposals vary on a single dimension extends to more general classes of heterogeneous preferences, and Proposition 5 sheds further light on what determines the key members as well as the reduced form bargaining solution that is required to describe the limit agreement set as members get patient.

3.4.3. Further comments.

Throughout the paper, we assume that all proposals are Pareto improvements over the status quo. In events where members have very heterogeneous preferences (think of large variations in bliss points), there may be no proposal that generates a payoff that all members would prefer over the status quo. In such settings, reducing the majority requirement could improve welfare (or even be Pareto-improving from an ex ante viewpoint), as it may allow to find out acceptable proposals (unlike in the unanimity case). We conjecture that such a consideration may have been a driving force behind the proposal of the Lisbon treaty (which includes the move from unanimity to qualified majority rules on a number of issues in the EU).

More generally, when we allow for the case in which $u_i(x)$ may be negative for some $x \in X$, or when we allow the utility functions $u_i(\cdot)$ not to be concave (they may still be assumed to be single-peaked), we still obtain the existence of two key members (who

would be indifferent between accepting or rejecting the highest possible proposal for one, and the smallest possible for the other) and we still obtain the existence of members who would be accepting all proposals (and for whom small changes in their bliss point would not affect the outcome). But, deriving a full characterization in such cases would be more difficult, as it would not necessarily be the case that the agreement set is an interval.²⁶

4. When proposals vary along more than one dimension.

In this Section, we consider situations in which proposals may affect members' preferences along more than one dimension. Specifically, we consider proposal spaces X having more than one dimension, and we assume that members care about the various dimensions of X (see the exact conditions in Proposition 7). Such a setup is well suited to deal with collective decisions bearing on several issues which may affect differently the various members.

As in Section 3, our analysis of the multi-dimensional proposal space case first focuses on whether the agreement set is small, as members get patient. While the agreement set remains small under the unanimity rule no matter what the number of dimensions of the proposal space is, the agreement set is typically large under any other majority requirement whenever the proposal space varies along sufficiently many dimensions. Thus, the distribution from which proposals are drawn may now have a significant impact, even as members are very patient.

We next turn to the analysis of what determines the locus of the agreement set when it is small. We first provide an example suggesting how the insight that members with moderate preferences do not affect the agreement set under unanimity carries over to proposal space cases with dimensions higher than one. Then we show that when the proposal space varies along as many dimensions as there are members, then all members affect the locus of the agreement set under the unanimity rule. We characterize the agreement set in this case.

²⁶In addition, uniqueness of equilibrium would no longer be guaranteed.

4.1. The size of the agreement set.

A general observation is that under unanimity, the acceptance set must be small as members get patient. For later use, we allow members to have different (yet not too dissimilar) discount factors. Formally, we define:

$$\eta \equiv \max_{i,j} \frac{1 - \delta_i}{1 - \delta_j} - 1.$$

η is a measure of the heterogeneity of members' patience. When members are equally patient, $\eta = 0$ and the more heterogeneous members are (in terms of patience) the larger η . The following result says that, for not too heterogeneous patiences, under unanimity, the agreement set gets small as members get patient.

Proposition 6: *Consider the unanimity rule. For any $\varepsilon > 0$, and any $\eta_0 > 0$, there exists δ_0 such that if $\delta_i \geq \delta_0$ for all i and $\eta \leq \eta_0$, then for any $f \in \mathcal{F}$, any equilibrium agreement set has size below ε .*

Intuitively, the agreement set cannot be so small that it generates costly delays to all, because otherwise, many proposals would become attractive to all, and the agreement set would thus not be small. Given that there cannot be delays that are costly to all, there must be a (patient) member, say i , that obtains a payoff that does not differ much from the expectation $\tilde{u}_i \equiv E[u_i(x) \mid x \in A]$. Since under unanimity, each member has the option to veto any proposal, all proposals in A must yield member i a payoff that cannot be much below \tilde{u}_i . Given that u_i is locally non-constant, this implies that A has a small size. We provide a formal proof in the Appendix.

We have shown in Section 3 that when X has only one dimension, qualified majority rules k with $k > n/2$ lead to small agreement sets, as members get very patient. We now point out that with majority rules other than unanimity this need not be so when the space of proposals vary along more than one dimension.

Proposition 7: *Consider the k -majority rule with $k < n$. If X is convex, has dimension $m > k$ and has a locally differentiable boundary, and if the gradient matrix of $u = (u_1, \dots, u_n)$ has rank at least $k + 1$ for every $x \in X$,*

then there exists $\rho > 0$ such that for all discount factors, the acceptance set has a size no less than ρ .

The proof of Proposition 7 appears in the Appendix. The rank condition means that local variations in X induces variations in all directions for at least $k + 1$ players. The intuition for the proof is as follows. Consider the expected agreement point $x^* = E[x | A]$. Since X is convex, $x \in X$, and since u is concave, x^* is unanimously accepted. Now, whether x^* lies on the frontier of X or not, the rank condition immediately implies that a non-negligible chunk of offers would be preferred to x^* by at least $k + 1$ members, thereby implying that the agreement set is large.

To understand the content of Proposition 7 more concretely, consider the special case in which $u_i(x) = x_i$ and x is drawn from the simplex $X = \{x, x_i \geq 0, 0 \leq \sum_i x_i \leq 1\}$ according to the uniform density.²⁷ In a symmetric equilibrium, each member obtains an expected payoff that cannot exceed $1/n$. Thus, the acceptance set should include the set $X^{(k)} = \{x \in X, x_i > \frac{1}{n} \text{ for at least } k \text{ members}\}$ whatever δ . Since $X^{(k)}$ has a size which is bounded away from 0 whatever $k < n$, we conclude that the agreement set must be large even as members are very patient.

As mentioned in Introduction, a large agreement set derived in the limit of patient members implies that the collective character of the search may induce extra randomness in the decision process that is not due to the impatience of members. This illustrates an important qualitative difference between individual and collective search.

Dependence on the density f .

Clearly, when the agreement set remains large for all discount factors as in the context of Proposition 7, the distribution $f(\cdot)$ from which proposals are drawn has a significant impact on the distribution of accepted proposals, even as members are very patient. But, when the agreement set is small as in the contexts of Propositions 2 and 6, equilibrium outcomes are not very sensitive to changes in the density f . We make this observation formal now.

²⁷Note that the frontier of X is not smooth. The argument only requires that part of the frontier is smooth: the part of the frontier which has the Pareto frontier of $u(X)$ as image.

For any $f \in \mathcal{F}$, $\zeta \geq 0$ and $\delta < 1$, define $A(\delta, \zeta, f)$ as the union of equilibrium agreement sets that may obtain when the density over proposals is f , discount factors δ_i all exceed δ and are not too dissimilar in the sense that $\eta \equiv \max_{i,j} \frac{1-\delta_i}{1-\delta_j} - 1 \leq \zeta$.

Definition 1. *A collective search problem is a small agreement set setting if there exists $\zeta > 0$ such that for all $\varepsilon > 0$ there exists $\delta < 1$ such that for all $f \in F$, size $A(\delta, \zeta, f) < \varepsilon$.*

In other words, in a small agreement set setting, as members get more and more patient, the agreement set gets smaller and smaller uniformly over all possible densities, and this remains true even when some small heterogeneity in the discount factors is allowed. The settings considered in Section 3 and Subsection 4.2.2 are small agreement set settings.

We have:

Proposition 8: *Consider a small agreement set setting. For every $\varepsilon > 0$, there exist $A \subseteq X$ with size $A \leq \varepsilon$ and δ_0 such that for all $f \in F$ and $\delta > \delta_0$, $A(\delta, 0, f) \subset A$.*

Proposition 8 says that in a small agreement set setting, a change in the density f cannot generate big changes in the equilibrium agreement set when members are equally patient.²⁸ We will make use of this result when analyzing situations in which members can influence the distribution of proposals (see Section 6).

4.2. Who drives the outcome?

Our analysis in the one-dimensional case shows that the outcome is determined by the preferences of at most two members under any majority rule ($k > \frac{n}{2}$), and that under stronger majority requirements, the preferences of members more "moderate" than these two key members play no role in determining the outcome.

²⁸Observe that Proposition 8 does not follow directly from Definition 1 because even though size $A(\delta, \zeta, f)$ is small for every f in a small agreement set setting, it might have been that $A(\delta, 0, f)$ varies significantly with f . The idea of the proof is to relate the changes of $A(\delta, 0, f)$ as one varies f to small changes of η in $A(\delta, \eta, f_0)$ for some fixed f_0 .

In the multi-dimensional case, the picture is somewhat different. Whenever the agreement set remains large even for very patient members, as in the context of Proposition 7 (for the k -majority rule whenever $m > k$), one should typically expect that all members affect the shape of the agreement set.²⁹ Whenever the agreement set gets small with patient players as under the unanimity rule (see Proposition 6), we illustrate through a simple example that it may still be the case that some members have no effect on the agreement set, thereby shedding light on how to extend the notion of moderate members in multidimensional settings. But, such a situation requires that proposals vary along fewer dimensions than there are players. When $m = n$, and the unanimity rule prevails, the preferences of all members affect the locus of the agreement, as we shall see.

4.2.1. When some members do not affect the agreement set

We consider the same environment as in Section 3 except that proposals now vary in two dimensions. Specifically, we assume that $X = [-1, 1]^2$, and that members have preferences of the form

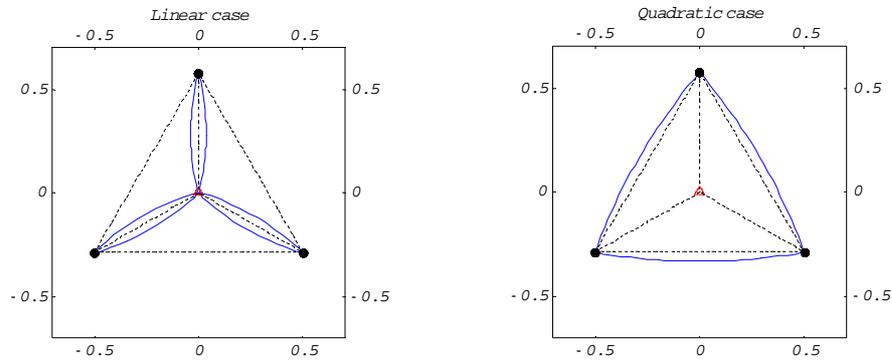
$$u_i(x) = v(x, \theta_i) \text{ with } v(x, \theta) \equiv 1 - \|x - \theta\|^a / 4,$$

where $\|x - \theta\|$ denotes the euclidean distance between x and θ , and where $a \geq 1$ is a parameter that reflects the risk aversion of members. In what follows, we focus on the linear case ($a = 1$) and the quadratic case ($a = 2$).

We shall also assume that bliss points of members are equidistant and at the same distance from the origin: $\|\theta_j - \theta_i\| = 1$ for $i, j \neq i$ and $\|\theta_i\| = \frac{1}{\sqrt{3}}$.

²⁹The reason is that with a large agreement set, each member will typically be pivotal on part of the boundary of the agreement set. To get a sense of this, assume that $m = n$, that members are arbitrarily patient, and that the agreement set is large. Then the equilibrium threshold vector \underline{u} coincides with $E[u(x) | A]$ and it lies in the interior of $u(X)$. Consider \underline{x} such that $u(\underline{x}) = \underline{u}$. If the gradient matrix of $u = (u_1, \dots, u_n)$ has rank n , then looking at a neighborhood of \underline{x} , one can check that the set of proposals x that are accepted by exactly $k - 1$ members other than i and $u_i(x) = \underline{u}$ is non empty, so any member i is pivotal on some boundaries of the agreement set, thereby implying that the agreement set is affected by the preferences of all members.

Consider the unanimity rule. When members are patient, the agreement set gets small and shrinks towards the origin $(0,0)$ (i.e. the small red set in the figures below). For a given discount factor δ close to 1, we are interested in whether adding another member would affect the agreement set, still assuming that the decision process is governed by the unanimity rule. This other member is assumed to have preferences also characterized by v , but with a possibly different bliss point θ . The following figures give the boundary of a set defined as follows: so long as the bliss point θ of the extra member lies within that set, the additional member has no effect on the agreement set.



Note that it is not surprising that a new member with bliss point located at $\theta_0 = (0,0)$ would have no effect on the agreement set. Indeed, such a member can get almost what he likes best by letting the other three agents decide on their own, so why would he veto any proposal jointly accepted by the other three members? Following the insight developed in the one-dimensional case, it can be shown that any new member with bliss point located on a segment $[\theta_0, \theta_i]$ would have no effect either. Indeed, due to the concavity of preferences, such a member is less eager to veto a proposal in the agreement set furthest away from his bliss point than member θ_i . What the figures show is that the set of members that would have no effect on the agreement set may actually be a much larger set. Whether this set is thin or thick depends on the concavity of the preferences: the larger a , the larger the set. For quadratic preferences, it includes all bliss points lying within the triangle delineated by θ_1 , θ_2 and θ_3 .³⁰

³⁰These figures have been drawn in two steps. First, restricting attention to the three members 1,2

4.2.2. The rich proposal space.

We assume now that the space of proposals is *rich* and that preferences are generic. That is, we assume that local variations in the space of proposals generate all possible variations in the utility space. So in particular, the dimension of the space of proposals (m) must be at least as large as the number of committee members (n).

Formally, we make the following assumption, which not only ensures that the space of proposals is rich, but also that the Nash bargaining solution among all members is uniquely defined and that it is a non-degenerate point of the Pareto frontier:

Assumption 3: *Assume that (i) $u : X \rightarrow u(X)$ has everywhere a gradient matrix with rank n (ii) $u(X) = (u_i(X))_{i=1}^n$ is a convex set with a locally differentiable boundary, and (iii) the generalized n -person Nash bargaining solution, $v^* = u(x^*)$ where $x^* = \arg \max_{x \in X} \prod_i u_i(x)$, is such that $u_i(x^*) > \min_{x \in X} u_i(x)$.*

One immediate corollary of Proposition 6 (agreement set is small under unanimity) is that, when members are very patient, the agreement set A must yield payoffs that lie close to the Pareto frontier of $u(X)$ under unanimity. Indeed, otherwise, all proposals that lead to a Pareto improvement over $E(u(x) \mid x \in A)$ would be unanimously accepted, hence the agreement set would be large. We show below that equilibrium outcomes must get close to the generalized n -person Nash bargaining solution,³¹ thereby implying that the preferences of all members affect the locus of the agreement set in this case.

and 3, we consider the agreement set that results when the discount δ is chosen close enough to 1 so that $\lambda = 9/10$. Having defined the agreement set (A), we find out numerically the set of points θ for which $u(x - \theta) > \lambda E[u(x - \theta) \mid A]$ for all $x \in A$.

³¹This can be viewed as the analog of Binmore et al. (1986) in our random offer bargaining setup. Note that we allow for more than two players, but yet restrict attention to stationary equilibria. Wilson (2001) obtains a similar characterization for the case of two players and he makes conjectures for the more than two player case.

Proposition 9: *Let Assumption 3 hold. For any density $f \in \mathcal{F}$, when δ tends to 1, equilibrium values tend to the generalized n -person Nash bargaining solution v^* .*

4.3. Discussion

Heterogeneous patiences: The insights developed in subsection 4.2 extend to the case of asymmetrically impatient members. Specifically, letting α_i be such that $(1 - \delta_i) = (1 - \delta)/\alpha_i$, one can still construct situations as in 4.2.1 in which some members do not affect the shape of the agreement set under the unanimity rule. Concerning Proposition 9, define now $v^{**} = u(x^{**})$ where $x^{**} = \arg \max_{u \in u(X)} \prod_i (u_i(x))^{\alpha_i}$ and assume that Assumption 3 holds with x^{**} (instead of x^*). We have that under the unanimity rule equilibrium values tend to v^{**} as δ tends to 1 (see the end of the proof of Proposition 9 in the Appendix).

Efficiency considerations: Comparing the efficiency of the various majority rules goes beyond the scope of this paper. Here, we make simple observations that are immediate corollaries of previous results. From Propositions 6 and 7, we can infer that when $u(X)$ is n -dimensional, the profile of payoffs is bounded away from the Pareto frontier under any k -majority rule other than the unanimity rule. Thus, for symmetric problems,³² the unanimity rule dominates any k -majority rule (with $k < n$) when members are patient enough and X varies along as many dimensions as there are members.

By contrast, consider the case in which $u_i(x) = x_i$ and x is drawn from the uniform distribution on $\{x, \sum_i x_i = 1\}$. It is then easily seen that the unanimity rule is dominated by any qualified majority rule given that the acceptance set is much smaller under unanimity than under other qualified majority rules and the equilibrium payoff obtained by every member is an increasing function of the probability that proposal x falls in the agreement set.

³² Assuming $X = [0, 1]^n$, $f(\cdot)$ is symmetric in x if $u_i(x) = u_j(x')$ whenever x and x' are obtained by permuting the i^{th} and j^{th} component.

These two observations suggest a trade-off between the unanimity rule and less demanding majority requirements. When the examined proposals are welfare equivalent, less demanding majority requirements are preferable to unanimity (because they speed up the agreement). When proposals are not welfare-equivalent, unanimity is preferable to other majority rules, as members get sufficiently patient. This trade-off is further discussed in Compte and Jehiel (2004) (see also Albrecht et al. (2009)).

5. Comparison with the random proposer model.

In the random proposer model (Binmore 1987), each party is selected with probability $1/n$ to make an offer. Starting with Baron and Ferejohn's seminal paper, this model has been used in numerous political science applications. We review how our model compares with the random proposer model.

Both models yield the same prediction (the Nash bargaining outcome) when players are patient, the set of proposals is rich and the decision rule is the unanimity (see Binmore et al. (1986) for the two-player case). They also yield the same prediction (median voter outcome) when the set of proposals is one-dimensional, preferences are single-peaked, players are patient, and the decision rule is the simple majority rule (see Banks and Duggan (2000) for the random proposer model).

Otherwise, the models generate different predictions, and we now review these differences. First, in the random proposer model, there is always a limited number of offers made in equilibrium (typically, just one per player). By contrast, the agreement set may have a positive measure even in the limit as members are very patient in our collective search model, when the dimension of proposal space is big enough and a qualified majority rule other than unanimity prevails.

Second, when proposals vary along a single dimension and majority rules other than simple majority are considered, our model predicts that at most two members determine the acceptance set whereas the random proposer model would predict that all members affect the decision. For concreteness, consider the unanimity case. Our search model predicts that only the preferences of the extremists matter. In the random proposer model, the solution would not coincide with that of our model, because the way bliss

points are distributed over the segment $[0, 1]$ would matter.³³ An equilibrium would consist of a pair $\{\underline{x}, \bar{x}\}$ of proposals: members with bliss point below \underline{x} would offer \underline{x} , and members with bliss point above \bar{x} would offer \bar{x} . The relative frequency with which \underline{x} and \bar{x} are proposed would thus depend on the number of members with bliss points below and above \underline{x} and \bar{x} , and so would the locus of \underline{x} and \bar{x} . As δ tends to 1, the solution would tend to the weighted Nash bargaining solution among the two most extreme members (θ_1 and θ_n), in which weights are determined endogenously by the distribution of members along the segment (θ_1, θ_n) .

6. Unifying the random proposer and the collective search models.

In this Section, we suggest a model that unifies the random proposer and the collective search models. Specifically, we define *a bargaining model with imperfect control*. We consider a setup in which at the start of every period, members simultaneously exert efforts to generate proposals: member i exerts effort $e_i \in [0, 1]$, at some cost $c_i(e_i)$. When the profile of effort is $e = (e_i)_{i=1}^n$, the proposal in this period is drawn from a density $f(\cdot | e)$ on X . After the proposal is drawn, every member i votes on whether he supports or not the proposal, and it is implemented whenever it receives at least k positive votes in the k -majority rule. We assume that all members are equally patient, $\delta_i = \delta$ for all i .

The imperfect control bargaining model encompasses the random proposer model of Baron and Ferejohn by letting $c_i(\cdot) \equiv 0$ and having $f(\cdot | e)$ put equal weight on the proposals chosen by each member upon being a proposer. It also encompasses the *collective search model* previously considered by letting $c_i(0) = 0$, $c_i(e_i) \equiv \infty$ for $e_i \neq 0$, and identifying $f(\cdot | 0)$ with the distribution over proposals previously considered, say $f(\cdot | 0) \equiv f_0(\cdot)$. But, it also covers many situations in between, for example allowing to capture the idea of lobbying activities and also the possibility that no proposal be found (simply by putting weight on the status quo in $f(\cdot | e)$).

For the next result, we consider the following assumption:

³³Note that the equilibrium value vector would not coincide with the generalized Nash solution either.

Assumption 4: (i) For every i , c_i is a smooth increasing function of e_i with $c_i(0) = 0$; (ii) $f(\cdot | e) \in \mathcal{F}$ for all $e \in [0, 1]^n$.

We have:

Proposition 10. *Let Assumption 4 hold and consider a small agreement set setting as defined in Definition 1. The equilibrium in the imperfect control bargaining model is such that as members get very patient only proposals very close to the agreement set in the collective search model are implemented and members' cost of effort gets very small.*

The argument is an immediate corollary of Proposition 8 which implies that in small agreement set settings with patient members, the set of possible equilibrium outcomes varies little with the distribution over proposals. Thus, under Assumption 4, the equilibrium outcome does not vary much with the effort profile, thereby implying that members cannot be willing to exert significantly costly effort in equilibrium.

Proposition 10 has various implications. First, if Assumption 4 holds and members are patient, we should expect to see more wasteful effort being exerted under qualified majority rules than under unanimity when the proposals vary along sufficiently many dimensions, while no such distinction should be expected when proposals vary along a single dimension. This is because when the space of proposal is rich enough and the decision rule is different from the unanimity rule, the agreement set is not small, and influencing the distribution of proposals may then have some value. So even when it is costly to affect the distribution of proposals, we should expect members to exert significant effort in equilibrium.³⁴

Second, Proposition 10 has interesting implications concerning the robustness of two classic insights in bargaining theory. The traditional view is that bargaining power is driven by the relative impatience (Rubinstein (1982)) and by the relative frequency

³⁴More generally, we should expect that more wasteful activities are exerted as the majority requirement is decreased since the agreement set gets larger. Yildirim (2007) makes a related observation in his model of influence activities. In the symmetric case, he shows that parties exert more effort to be the proposer as the majority requirement is less strong (see his Proposition 5).

with which parties make offers. While the effect of relative impatience is robust to the introduction of imperfect control of offers, the frequency with which parties make offers plays little role when parties are sufficiently patient and only imperfectly control the offers being made (Proposition 10).³⁵

Finally, we note that when Assumption 4 does not hold - for example, because when members exert too little effort only Pareto inferior proposals are generated, and when agreement sets are small in the collective search model, there is a possibility of under-provision of effort in the imperfect control bargaining model. If the proposal space is rich enough, this may lead to prefer qualified majority to unanimity, so as to induce larger acceptance sets hence stronger incentives to provide effort.

7. Conclusion

This paper has introduced a model of collective decision process in which parties lack control over the proposals put to a vote, which represents well those situations in which new projects must be adopted and some exogenous forces (such as the generation of ideas) affect the projects to be considered. Our results bear on the effect of this lack of control as parties get very patient and on who has most impact on the implemented decisions.

Two sorts of results are obtained and they contrast with those that would obtain if parties had perfect control. First, when proposals vary along a single dimension, only two members determine the distribution of accepted proposals whereas all members play a role in the perfect control case. Second, the acceptance set may remain large under qualified majority rules other than unanimity when proposals vary along sufficiently many dimensions whereas by contrast at most as many offers as there are players can be made in the perfect control case.

Unifying our collective search model and the random proposer model has revealed the robustness of the insights derived in the extreme case in which members have no control over the proposals put to a vote, and it has suggested new avenues for the study

³⁵Which prediction applies (the search model or standard bargaining) depends on the order of the limits, that is, how the noisiness of the proposal process compares to patience.

of lobbying activities in collective bargaining.

Even though participation is left exogenous in our model, some of the insights we derive can be related to the results obtained in the political economy literature that analyze who has most incentives to participate in collective decision processes. Osborne et al. (2000) obtain for a number of specifications of how the profile of participants affects the decision outcome (these are left exogenous in Osborne et al) that the most extremists have greater incentives to participate in the decision process. Our insight that under unanimity the extremists determine the final decision when proposals vary along a single dimension would also lead to the conclusion that extremists have greater incentives to participate in the decision process under that rule. More generally, our analysis, which sheds light on how the rules of the decision process affect the final outcome as a function of the participants, can be used to analyze the incentives of the various members to participate.

References

- [1] Albrecht, J., A. Anderson and S. Vroman (2009): 'Search by Committee,' unpublished
- [2] Banks, J. and Duggan, J. (2000) "A Bargaining Model of Collective Choice", *American Political Science Review*, **94**, 73-88.
- [3] Baron, D. and Ferejohn (1989), "Bargaining in Legislatures", *American Political Science Review*, **90**, 316-330.
- [4] Binmore, K. (1987) "Perfect Equilibria in Bargaining Models", in *The Economics of Bargaining*, ed. K. Binmore and P. Dasgupta, Oxford: Basil Blackwell, 77-105.
- [5] Binmore, K., A. Rubinstein and A. Wolinsky (1986): "The Nash Bargaining Solution", *RAND Journal of Economics*, **17**, 176-188.
- [6] Compte, O. and P. Jehiel (2004): "Bargaining over randomly generated offers," mimeo

- [7] Downs, A. (1957), *An Economic Theory of Democracy*, Harper and Row, N.Y., 1957.
- [8] McCall, J. (1970): 'Economics of Information and Job Search,' *Quarterly Journal of Economics*, **84**, 113-126.
- [9] Nash, J. (1950), "The Bargaining Problem," *Econometrica* **18**, 155-162.
- [10] Osborne, M. J. Rosenthal and M. Turner (2000): 'Meetings with Costly Participation,' *American Economic Review*, **90**, 927-943.
- [11] Rogerson R., R. Shimer and R. Wright (2005): 'Search-theoretic Models of the Labor Market: a Survey,' *Journal of Economic Literature*, **43**, 959-988.
- [12] Rubinstein, A. (1982): "Perfect Equilibrium in a Bargaining Model", *Econometrica* **50**, 97-109.
- [13] Strulovici, B. (2007): "Voting and Experimentation," mimeo.
- [14] Wilson, C. (2001): "Mediation and the Nash bargaining solution," *Review of Economic Design*, **6**, 353-370.
- [15] Yildirim, H. (2007): "Proposal Power and Majority Rule in Multilateral Bargaining with Costly Recognition," *Journal of Economic Theory* **136**, 167-196.

Appendix

Proof of Proposition 2. Consider an equilibrium. Let A denote the agreement set, and denote respectively by \bar{x} and \underline{x} the supremum and infimum value in A . The acceptance threshold for player i writes as

$$\underline{u}_i = \delta v_i = \lambda E[u_i(x) \mid x \in A]$$

where $\lambda = \frac{\delta \Pr A}{1 - \delta + \delta \Pr A}$. Let \tilde{x} denote the expected agreement point, that is, $\tilde{x} = E[x \mid x \in A]$. Since u_i is concave, and since u_i is positive and $\lambda < 1$, we have (applying Jensen's inequality):

$$u_i(\tilde{x}) \geq E[u_i(x) \mid x \in A] > \lambda E[u_i(x) \mid x \in A] = \underline{u}_i$$

It follows that \tilde{x} is unanimously accepted. Now observe that if $u_i(\bar{x}) \geq \underline{u}_i$, then, since u_i is single peaked and $u_i(\tilde{x}) > \underline{u}_i$, we must also have $u_i(x) > \underline{u}_i$ for all $x \in [\tilde{x}, \bar{x}]$. So if there is a qualified majority k for \bar{x} , there must also be a qualified majority for any $x \in [\tilde{x}, \bar{x}]$. The same argument applies to the interval $(\underline{x}, \tilde{x}]$, implying that A is an interval.

Now fix any $\varepsilon > 0$ and assume that agreement set has size at least equal to ε , that is, $\bar{x} - \underline{x} \geq \varepsilon$. Because preferences are not locally constant (see footnote 6) and $f(\cdot)$ is bounded away from 0, there exists a constant a such that for any interval A of size at least ε , and for every member i ,

$$E[u_i(x) \mid x \in A] \geq \min\{u_i(\underline{x}), u_i(\bar{x})\} + a$$

Since the distribution over draws puts weight on all values of x , $\Pr A$ has a lower bound, so for δ close enough to 1, λ is arbitrarily close to 1, hence player i must either reject draws close to \underline{x} or reject draws close to \bar{x} . It follows that when $k > n/2$, if there are k players that accept draws close to \underline{x} , these k players must reject draws close to \bar{x} , contradicting the premise that $\bar{x} \in A$. **Q.E.D.**

Proof of Proposition 3: Define $g(\theta, x) = v(x - \theta) - \lambda E[v(x - \theta) \mid A]$. We shall use the following observation, which will be proved later on.

Lemma 1: For any θ : if $\theta < \bar{x}$ then $\frac{\partial g}{\partial \theta}(\theta, \bar{x}) > 0$; if $\theta > \underline{x}$ then $\frac{\partial g}{\partial \theta}(\theta, \underline{x}) < 0$; if $\theta \geq \bar{x}$, then $g(\theta, \bar{x}) > 0$; if $\theta < \underline{x}$ then $g(\theta, \underline{x}) > 0$.

(i) There must exist i_0 such that $g(\theta_{i_0}, \bar{x}) = 0$ because otherwise, either $u_i(\bar{x}) > \delta v_i$ for at least k members, and by continuity, there are proposals $x > \bar{x}$ that would be accepted as well. Or $u_i(\bar{x}) < \delta v_i$ for at least $n - k$ members, and all proposals in some small neighborhood of \bar{x} must be rejected as well, contradicting the premise that $\bar{x} \in A$. Similarly, there must exist i_1 such that $g(\theta_{i_1}, \underline{x}) = 0$.

(ii) Assume $g(\theta_{i_0}, \bar{x}) = 0$. Then Lemma 1 implies that $\theta_{i_0} < \bar{x}$ (as otherwise i_0 would not be pivotal for x close to \bar{x} and the same argument as in (i) could be applied to get a contradiction). It follows from Lemma 1 that for all $\theta < \theta_{i_0}$, $g(\bar{x}, \theta) < 0$ and that for all $\theta > \theta_{i_0}$, $g(\bar{x}, \theta) > 0$.³⁶ Equivalently, for all $i < i_0$, $u_i(\bar{x}) < \delta v_i$, and for all $i > i_0$, $u_i(\bar{x}) > \delta v_i$. A similar argument shows that for all $i > i_1$, $u_i(\underline{x}) < \delta v_i$ and that for all $i < i_1$, $u_i(\underline{x}) > \delta v_i$.

(iii) Step (ii) implies that in order to have at least k players in favor of \bar{x}^- , one must have $i_0 \leq n - k + 1$, and that in order to have fewer than k players in favor of \bar{x}^+ , one must have $i_0 \geq n - k + 1$. Thus $i_0 = n - k + 1$.

A similar argument permits to show that $i_1 = k$.

(iv) Since $k > n/2$, we have $i_0 \leq i_1$. Consider now any two members $i_0 \leq i_1$ and the acceptance interval $A = [\underline{x}, \bar{x}]$ which satisfies $g(\theta_{i_0}, \bar{x}) = 0$ and $g(\theta_{i_1}, \underline{x}) = 0$. By step (ii), A is an equilibrium agreement set of the game with all players and qualified majority rule k . Consider the game where only i_0 and i_1 are present and unanimity is required. Since $i_0 \leq i_1$, then by step (ii), any $x \in A$ is accepted by both i_0 and i_1 , and any $x \notin A$ is rejected by either i_0 and i_1 , so A is an equilibrium agreement set of that game as well. **Q.E.D.**

Proof of Lemma 1: Since v is concave, for any $\theta < \bar{x}$, we have

$$\begin{aligned} \frac{\partial g}{\partial \theta}(\theta, \bar{x}) &= -v'(\bar{x} - \theta) + \lambda E[v'(x - \theta) \mid x \in A] \\ &\geq -(1 - \lambda)v'(\bar{x} - \theta) > 0. \end{aligned}$$

³⁶This is because either $\theta < \bar{x}$ and then $\frac{\partial g}{\partial \theta}(\theta, \bar{x}) > 0$, or $\theta > \bar{x}$ (and then θ must prefer \bar{x} to a draw from A obtained with delay).

and clearly for $\theta \geq \bar{x}$, $g(\theta, \bar{x}) > 0$ given that $\lambda < 1$. The same argument is used to prove that if $\theta > \underline{x}$ then $\frac{\partial g}{\partial \theta}(\theta, \underline{x}) < 0$ and if $\theta < \underline{x}$ then $g(\theta, \underline{x}) > 0$. **Q. E. D.**

Proof of Proposition 4: For the simple majority rule with an odd number of players, we already know that $k = 1 + n - k$ and that the proposal θ_k is always accepted. The result thus follows from the fact that A is small when δ tends to 1. In what follows, we restrict attention to the other cases and set $i_0 = 1 + n - k$ and $i_1 = k$. Proposition 3 shows that the agreement set $A = [\underline{x}, \bar{x}]$ satisfies three equations:

$$\begin{aligned}\lambda &= \frac{\delta \Pr A}{1 - \delta + \delta \Pr A} \\ u_{i_0}(\bar{x}) &= \lambda E[u_{i_0}(x) \mid x \in A] \\ u_{i_1}(\underline{x}) &= \lambda E[u_{i_1}(x) \mid x \in A]\end{aligned}$$

When δ tends to 1, A has a small size, say $\bar{x} - \underline{x} = \varepsilon$. We shall use the following approximation:

$$u_i(x) = u_i(\underline{x}) + (x - \underline{x})u'_i(\underline{x}) + O(\varepsilon^2).$$

Now consider the conditional distribution over draws $h(x) = \frac{f(x)}{\int_{\underline{x}}^x f(x)dx}$. Because f is bounded below and $|f'|$ bounded above, there exists a constant a such that for all $x \in A$, $|h(x) - 1| \leq a\varepsilon$. It follows that³⁷

$$E[u_i(x) \mid x \in A] = u_i(\underline{x}) + \frac{\varepsilon}{2}u'_i(\underline{x}) + O(\varepsilon^2).$$

Solving for λ and omitting terms of order 2 or larger in ε , we obtain:

$$(u_{i_0}(\underline{x}) + \varepsilon u'_{i_0}(\underline{x}))(u_{i_1}(\underline{x}) + \frac{\varepsilon}{2}u'_{i_1}(\underline{x})) = u_{i_1}(\underline{x})(u_{i_0}(\underline{x}) + \frac{\varepsilon}{2}u'_{i_0}(\underline{x})),$$

which implies that

$$u'_{i_1}(\underline{x})u_{i_0}(\underline{x}) + u_{i_1}(\underline{x})u'_{i_0}(\underline{x}) = O(\varepsilon^2).$$

Now let $h(x) = u'_{i_1}(x)u_{i_0}(x) + u_{i_1}(x)u'_{i_0}(x)$. The Nash bargaining solution $x^* = x_{i_0, i_1}^*$ solves $h(x^*) = 0$. Since h is continuous, and since h is strictly monotone on $(\theta_{i_0}, \theta_{i_1})$, \underline{x} must tend to x^* as ε gets small. **Q.E.D.**

³⁷Thus the dependence on the distribution over draws only appear in the term $O(\varepsilon^2)$, which explains why it will not affect the locus of the limit agreement set.

Proof of Proposition 5: Consider δ close to 1 and assume that i_0 and i_1 are the members referred to in Proposition 5. Consider the agreement set $A = [\underline{x}, \bar{x}]$ that obtains in the game where only i_0 and i_1 are present and unanimity prevails. By an argument identical to that of Proposition 4, the agreement set must be close to the solution of the equation $\mu_{i_0}(x) = \mu_{i_1}(x)$.

Now set $m = 1 + n - k$ and define $i_0^{(m)}$ (respectively $i_1^{(m)}$) as a member for which $\theta_i > x$ (respectively $\theta_i < x$) and $\mu_i(x^{(m)}) = \mu^+(x^{(m)})$. Consider any choice of i_0 and i_1 that would *not* satisfy these properties. By definition of $x^{(m)}$, and given the choice of m , one of the following properties must fail: there is not a k -majority for \bar{x}^+ or \underline{x}^- ; there is a k -majority for \bar{x}^- and \underline{x}^+ . **Q.E.D.**

Proof of Proposition 6: Consider an equilibrium. Let A denote the agreement set and let \tilde{x} denote the expected agreement point, that is, $\tilde{x} = E[x \mid x \in A]$. Recall that the acceptance threshold for player i writes as

$$\underline{u}_i = \delta_i v_i = \lambda_i E[u_i(x) \mid x \in A]$$

where $\lambda_i = \frac{\delta_i \Pr A}{1 - \delta_i + \delta_i \Pr A}$. Let $\bar{\lambda} = \max \lambda_i$ and $\bar{\delta} = \max \delta_i$. We first show that $\bar{\lambda}$ must get close to 1 when $\bar{\delta}$ tends to 1.

Since u_i is concave, $u_i(\tilde{x}) \geq E[u_i(x) \mid x \in A]$. Since u_i is continuously differentiable, there exists a such that for any ε , and for any x such that $\|x - \tilde{x}\| \leq \varepsilon$,

$$u_i(x) \geq u_i(\tilde{x}) - a\varepsilon$$

Choosing $\varepsilon \leq \frac{1 - \bar{\lambda}}{a} \min_{x \in X} u_i(x)$ thus ensures that for all i , $u_i(x) \geq \lambda_i u_i(\tilde{x}) \geq \underline{u}_i$, thereby implying that all x such that $\|x - \tilde{x}\| \leq \varepsilon$ are accepted. It follows that $\Pr A \geq b(1 - \bar{\lambda})$ for some constant b independent of $f \in \mathcal{F}$. Hence we have:

$$1 - \bar{\lambda} = \frac{1 - \bar{\delta}}{1 - \bar{\delta} + \bar{\delta} \Pr A} \leq \frac{1 - \bar{\delta}}{b(1 - \bar{\lambda})},$$

which implies $\bar{\lambda} \geq 1 - ((1 - \bar{\delta})/b)^{1/2}$.

Now fix $\varepsilon > 0$. Since u is locally non constant (see footnote 6), there exist a constant c such that for any convex set A of size larger than ε , and for any $f \in F$,

$$E[u_i(x) \mid x \in A] \geq \min_{x \in A} u_i(x) + c. \tag{7.1}$$

for at least one member i . Now assume that the agreement set A indeed has size larger than ε . The agreement set A is convex³⁸, so inequality (7.1) applies for some i . For that member i we have

$$\lambda_i E[u_i(x) \mid x \in A] \geq \min_{x \in A} u_i(x) + \lambda_i c - (1 - \lambda_i) \min_{x \in A} u_i(x).$$

Since $\eta \leq \eta_0$, λ_i tends to 1 when $\bar{\delta}$ tends to 1, hence player i does not want to accept all proposals in A when $\bar{\delta}$ is close to 1. **Q.E.D.**

Proof of Proposition 7: Consider $x^* = E[x \mid x \in A]$, where A is the equilibrium agreement set. X is convex so $x^* \in X$, and the functions u are concave, so x^* is unanimously accepted. Let $u_i^* = u_i(x^*)$. We wish to show that there exists $K \subset \{1, \dots, n\}$ with $|K| = k$ such that

$$\bigcap_{i \in K} \{x, u_i(x) > u_i^*\} \neq \emptyset. \quad (7.2)$$

Since the functions u_i are C_1 and since $f \in \mathcal{F}$, this will ensure that there is a subset X of measure bounded away from 0 that at least k members strictly prefer to x^* , which will thus ensure that A has measure bounded away from 0.

To show (7.2), define the half space $H_i = \{x \in R^m, \nabla u_i(x^*) \cdot (x - x^*) > 0\}$. If x^* lies in the interior of X , the rank condition ensures that one can find K with $|K| = k$ such that $\bigcap_{i \in K} H_i \neq \emptyset$. If x^* lies on the frontier of X , let g denote the normal and define $G = \{x, -g \cdot (x - x^*) > 0\}$. The rank condition ensures that one can find a subset \bar{K} of $k + 1$ linearly independent vectors $\nabla u_i(x^*)$. Choose $j, j' \in \bar{K}$. The rank condition ensures that if g is linearly dependant of $(\nabla u_i(x^*))_{i \in \bar{K} - \{j\}}$, it must be linearly independent of $(\nabla u_i(x^*))_{i \in \bar{K} - \{j'\}}$. It follows that one can find K with $|K| = k$ such that $G \cap \bigcap_{i \in K} H_i \neq \emptyset$. Since X is convex and smooth, and since the functions u_i are C_1 , we conclude that (7.2) holds. **Q. E. D.**

Proof of Proposition 8: Proposition 8 follows from Lemma 2.

³⁸For every $(v_i)_i$, $\bigcap_i \{x \mid u_i(x) > \delta v_i\}$ is the intersection of convex sets (thus convex itself) due to the concavity of u_i .

Lemma 2: *There exist two constant $\bar{\alpha}$ and b such that for any $f_0, f \in F$, for any $\delta_0 < 1$ and $\delta \geq 1 - b(1 - \delta_0)$, and for any $\varepsilon > 0$, if $\text{size}(A(\delta, 0, f)) \leq \varepsilon$, then $A(\delta, 0, f) \subset A(\delta_0, \bar{\alpha}\varepsilon, f_0)$.*

Lemma 2 relates the effect of a change of f on $A(\delta, 0, f)$ to the effect of a change of $\eta = \max \frac{1-\delta_i}{1-\delta_j} - 1$ on $A(\delta, \nu, f_0)$. To see how to apply Lemma 2, consider a small agreement set setting and fix ζ as in definition 1. Fix any $\varepsilon > 0$ such that $\varepsilon < \zeta/\bar{\alpha}$ and also fix some $f_0 \in F$. For δ_0 close enough to 1 and any $\delta \geq 1 - b(1 - \delta_0)$, by definition of a small agreement set setting, we have $\text{size}(A(\delta, 0, f)) \leq \varepsilon$ (which implies that Lemma 2 applies) and moreover $\text{size}(A(\delta_0, \zeta, f_0)) \leq \varepsilon$. Thus $A(\delta, 0, f) \subset A(\delta_0, \bar{\alpha}\varepsilon, f_0) \subset A(\delta_0, \zeta, f_0)$ (by Lemma 2), and Proposition 8 holds with $A \equiv A(\delta_0, \zeta, f_0)$.

Q. E. D.

Proof of Lemma 2: We refer to $\nu(\delta, \zeta, f)$ as the size of $A(\delta, \zeta, f)$. We choose b such that for any A and $f, f_0 \in F$, $\frac{\Pr_f(A)}{\Pr_{f_0}(A)} \geq 2b$. Consider now any $\delta_0, \delta \geq 1 - b(1 - \delta_0)$ and $f, f_0 \in F$. Fix $\varepsilon > 0$ and assume $\nu(\delta, 0, f) < \varepsilon$. At (δ, f) , consider an equilibrium with values $v = (v_i)_{i \in I}$ and with equilibrium agreement set A_f . We show below that $A_f \subset A(\delta, \bar{\alpha}\varepsilon, f_0)$.

Set $\underline{u}_i = \delta_i v_i$. Equilibrium conditions are:

$$\underline{u}_i = \frac{\Pr_f(A_f)}{1 - \delta} E_f[u_i(x) - \underline{u}_i \mid x \in A_f].$$

For any $f, f_0 \in \mathcal{F}$, if $\nu(\delta, 0, f) < \varepsilon$ then

$$E_f[u_i(x) - \underline{u}_i \mid x \in A] = E_{f_0}[u_i(x) - \underline{u}_i \mid x \in A](1 + \alpha_i \varepsilon)$$

for some $|\alpha_i| \leq \bar{\alpha}$, where $\bar{\alpha}$ is set independently of f . In addition, since $1 - \delta \leq b(1 - \delta_0)$ and since $\Pr_f(A_f) \geq 2b \Pr_{f_0}(A_f)$,

$$\frac{\Pr_{f_0}(A_f)}{1 - \delta_0} \leq \frac{1}{2} \frac{\Pr_f(A_f)}{1 - \delta}$$

One may thus find a vector of discount factors $(\delta_i)_{i \in I}$ where $\delta_i \geq \delta_0$ and $\frac{1-\delta_i}{1-\delta_j} \leq 1 + \bar{\alpha}\varepsilon$ for all i, j such that:

$$\underline{u}_i = \frac{\Pr_{f_0}(A_f)}{1 - \delta_i} E_{f_0}[u_i(x) - \underline{u}_i \mid x \in A_f].$$

thus implying that with discounts $(\delta_i)_{i \in I}$ and density f_0 , A_f is an equilibrium agreement set. **Q. E. D.**

Proof of Proposition 9: In what follows, $g(u) \equiv d(u, P)$ denotes the Euclidean distance from u to the Pareto frontier P . So $g(u) = 0$ corresponds to a parameterization of the Pareto Frontier.

Consider the distribution over utility profiles (u_1, \dots, u_n) over $u(X)$ induced by $f(\cdot)$. Since $f(\cdot)$ is bounded away from 0 and continuously differentiable on X , and since $u : X \rightarrow u(X)$ is a diffeomorphism, that distribution admits a density $h(\cdot)$ that is also bounded away from 0 on $u(X)$ and continuously differentiable.

Consider now $\varepsilon > 0$. There exists a scalar a such that for any set A of size ε , the conditional distribution on $u(A)$, satisfies:

$$|h(u) - 1| \leq a\varepsilon$$

Consider now an equilibrium value profile v , and define $\underline{u}_i = \delta v_i$. We are trying to characterize the vector $\underline{u} = (\underline{u}_1, \dots, \underline{u}_n)$ knowing $g(\underline{u}) \leq \varepsilon$. For each possible \underline{u} , the agreement set in the utility space is defined as:

$$D_{\underline{u}} = \{u \in u(X), u_i \geq \underline{u}_i \text{ for all } i\}$$

For each \underline{u} and member i , let $\bar{u}_i(\underline{u})$ denote the highest value that player i obtains in $D_{\underline{u}}$, and define the simplex

$$\bar{D}_{\underline{u}} = \{u \in u(X), u_i \geq \underline{u}_i \text{ for all } i, \sum_i \frac{u_i - \underline{u}_i}{\bar{u}_i(\underline{u}) - \underline{u}_i} \leq 1\}$$

Since the set $u(X)$ is convex with a smooth boundary, for \underline{u} such that $g(\underline{u}) \leq \varepsilon$, D contains \bar{D} and $D - \bar{D}$ has a measure comparable to ε^2 at most. Since $h(u)$ differs from 1 by at most $a\varepsilon$, the equilibrium condition (2.2) implies

$$\frac{\underline{u}_i}{\underline{u}_1} = \frac{E_h(u_i - \underline{u}_i \mid D)}{E_h(u_1 - \underline{u}_1 \mid D)} = \frac{\bar{u}_i(\underline{u}) - \underline{u}_i}{\bar{u}_1(\underline{u}) - \underline{u}_1} (1 + O(\varepsilon)) \quad (7.3)$$

Recall now that the equation $g(u) = 0$ corresponds to a parameterization of the frontier. Consider u such that $g(u) \leq \varepsilon$. We have:

$$0 = g(\bar{u}_i(u), u_{-i}) = g(u) + g'_i(u)(\bar{u}_i(u) - u_i) + O(\varepsilon^2).$$

The equilibrium threshold vector must thus have the property that for all i, j ,

$$\frac{u_j g'_j(u)}{u_i g'_i(u)} = 1 + O(\varepsilon) \quad (7.4)$$

Now the connection with the Nash bargaining is as follows. Consider a vector u such that $g(u) = \varepsilon$ and define $\beta_i = u_i g'_i(u)$. The vector u can be viewed as the (unique) point that maximizes $\prod u_i^{\beta_i}$ over the set $g(u) = \varepsilon$.

We know from Proposition 6 that \underline{u} must tend to the Pareto Frontier. When δ tends to 1, (7.4) shows that the weights β_i must get close to one another, hence \underline{u} must get close to the (unique) point that maximizes $\prod u_i$ over the set $g(u) = 0$.

Comment: To see the connection with the generalized Nash bargaining solution when discount factors differ, assume $1 - \delta_i = (1 - \delta)/\alpha_i$. Then the right hand side of (7.3) becomes $\frac{\alpha_i \underline{u}_i}{\alpha_1 \underline{u}_1}$, which then implies that (7.4) becomes $\frac{u_j g'_j(u)}{u_i g'_i(u)} = \frac{\alpha_i}{\alpha_j} + O(\varepsilon)$, which further implies that the weights β_i must get close to α_i (up to a multiplicative constant). **Q. E. D.**