Assumptions and Ommissions:

An Overview and Evaluation of the Banzhaf Power Index as a Tool for Assessing Voter Power in a Weighted Voting System

April 12, 2007

Ronan Geoghegan
ID 88239041
ISCI 330

In many governments and international organizations throughout the world a series of ongoing debates are raging, and while this in itself is not all that surprising, what might be is that the point of contention in many of these debates is one of the most fundamental of institutional processes; how decisions are made. While this is an issue that arises whenever a collection of individuals needs to come together to make a decision as a group, it is nevertheless very hard to resolve and there exists a variety of concerns that need to be addressed prior to doing so. For example, should each member have an equal say, or vote, in the decision making process or should some have more and others less? Furthermore, if some individuals are to have a greater say than others, how do we determine who and by how much? While game theory can not provide the answer to these questions, it can help to determine whether the decision making process that is developed as a result is fair or not. The Banzhaf Power Index (BPI) is a tool that is often employed to do just that, and is applied to assess the relative decision-making power of each individual actor when those actors are present in a system of weighted voting. That is, in a system in which the number of votes assigned to each actor is not equal but instead varies according to some external characteristic, such as population or financial contribution to the group. However, while this method does reveal some underlying truths regarding the theoretical distribution of power within a group, it is important to realize its limitations as a measure of power. Indeed, upon closer inspection it becomes clear that the BPI alone often can not fully reflect the actual distribution of voting power in a group. Three of the main reasons for this are that first the BPI assumes all coalitions in a weighted voting system are of equal value and thereby it disregards an important aspect of power known as P-power, second that it assumes all coalitions are equally likely to occur and thereby it disregards any player preferences, and finally that as it is a measure of relative power, the BPI fails to reveal any change in absolute power should the decision making body or decision rule be modified in any way. Nevertheless, before getting in to the limitations of the index itself, it is first necessary to outline what exactly the BPI is and how it is calculated.

The BPI is a construct of cooperative game theory, and is based on an *n*-person simple game in which each player has a different voting weight, and a certain quota of votes is needed for a decision, or resolution, to pass. As a result, the different players are forced to cooperate with one another to form a variety of coalitions, each of which may or may not have the required amount of votes necessary to pass a resolution. If a coalition does possess enough votes to pass the resolution, then it is considered a winning coalition. Moreover, any member of that coalition that can turn it from a winning one to a losing one by defecting from the coalition is considered a critical member. For example, in a group consisting of three players, *a*, *b*, and *c*, where players *a* and *b* have two votes and player *c* only has one, and in which the quota needed to pass a resolution is three out of the five votes, then any combination of two players will comprise a winning coalition. That is, a coalition between players *a* and *c* or between players *b* and *c* would possess three votes while that between *a* and *b* would possess four, so therefore all three are winning combinations. Furthermore, in all three of these two-player coalitions, both players are critical members as their defection would reduce the number of votes to below three, while in a three-player winning coalition no player is critical. Which is where the BPI comes in; the BPI can be defined as “the number of winning coalitions in which the member's defection from the coalition would render it losing … divided by the total number of critical defections for all players."[1] In this case each player is a critical member twice and the total number of critical defections is six, so each player has a BPI value of 1/3, resulting in a power vector of (1/3, 1/3, 1/3). So while at first glance it may have seemed as though player *c* had less power as a consequence of having half as much voting weight as the other players, it actually exerts as much influence over the decision as the other two as all three players can join the same number of winning coalitions. In this sense then, the voter weight often does not fully equate to voter strength and it is the aim of the BPI to assess what true underlying power each actor actually has.

The first limitation arises however, in just how the BPI defines the concept of power. According to Banzhaf himself, power can be defined as "the ability to affect outcomes," such as by defecting from a coalition and thereby changing the outcome of a vote from a pass to a fail.[2] This type of power has been termed I-Power, for the ability to influence. Yet another type of power, P-power, is not considered by the BPI at all as the BPI assumes that all coalitions are of equal value, which is an incorrect assumption. P-power is power defined as a voter's expected relative share in some prize that is only available to members of the winning coalition, and which varies between different coalitions.[3] For example, if a coalition has two members and both are critical, then they will share the 'profits' half and half, but if the coalition has three members and all are critical, then they will divide the profits among all three players thereby receiving a

[1] Flanagan, p 95.
[2] Braham & Holler, p. 138.
[3] Felsenthal & Machover (2001), p. 84.

smaller share than those in the two-member coalition.[4] This can be interpreted in real life situations by the fact that to form a successful three-member coalition, each member would likely have to make more compromises than if they only needed to cooperate with one other actor, and therefore the three-member coalition would be harder to form and result in a smaller payoff. In any event, this aspect of power is not considered by the BPI and consequently any conclusions it makes on the relative measure of power within a group would be different if this second aspect of power was taken in to consideration.

This is exactly what another index of power, the Johnston Power Index (JPI), strives to accomplish. The JPI is very much like the BPI, but differs in how it scores a player's criticalness and in its attempt to distinguish between different coalitions, which it does by giving a premium to actors in a coalition that are uniquely critical. While the BPI assigns a value of one each time a member is critical, the JPI takes in to account this 'unique criticalness' by assigning a value of $1/n$ points to a critical member in a winning coalition, where n is the number of total critical members in that coalition. The JPI, then, is the sum of points from a member's critical coalitions divided by the sum of all players' critical coalition points. While the two indices differ only slightly in their definitions, the differences this can have on the evaluation of power in a group is striking. According to a survey of the voting power in the Executive Board of the International Monetary Fund, Jonathan Strand in *International Interactions* concluded that the United States holds 63% of the voting power if the JPI is applied, but only about 27.1% if the BPI is.[5] By disregarding P-power in its definition of voting power and assuming all coalitions are of

[4] Felsenthal & Machover (2005), p 501.
[5] Strand, p. 29.

equal value, the BPI has ignored a significant aspect of power and so can only be considered partially representative of the real distribution of power within a group

The second false assumption that the BPI makes regarding coalitions is that it assumes not only that they are all of equal value, but also that every mathematically possible coalition is equally likely to occur. Yet this assumption does not coincide with real-life situations of collective decision making as it disregards player preferences and therefore misses an integral component of social power; strategic interaction. For example, in an international organization such as the United Nations, it is unrealistic to assume that every possible combination of votes is equally likely to occur but instead it is far more likely that countries in similar regions or of similar economic status would tend to vote together while those in other regions or of other status would tend to vote in a different way. This was famously witnessed during the Cold War when the Soviet Union and the United States, along with their respective blocs, would seldom vote together on any issue, thus bringing the United Nations to a standstill. So while it was possible for them to vote together and pass a resolution, that never actually occurred. Yet such a strategic interplay is missed by the BPI as it assumes all coalitions are equally probable, and so the BPI can at best describe *a priori* relative voting power, the power that exists before any player preferences develop and before any bargaining takes place. The BPI therefore can lay the foundations of relative power, but it can not take in to account the interplay that occurs beyond that level and that may also affect the power distribution of individuals in the group.[6]

---

[6] One method which tries to take in to account player preferences in weighted voting situations is that developed by Steuneneberg, Schmidtchen, and Kobalt. In this method, different states, such as the status

Finally, a third limitation of the BPI is that even if we do accept the assumptions that all coalitions are of equal value and are of equal likelihood to develop, then it still can not paint a complete picture by itself, but must rather be supplemented by other sets of data. Specifically, while it can provide a measure of relative power, the BPI fails to reveal any change in absolute power if the decision making rule of a group is modified. In *Game Theory and Canadian Politics*, Thomas Flanagan illustrates this fact with an example concerning the Canadian government. During the 1995 Quebec Referendum, one concession that Prime Minister Chretien made to appease the separatists was to promise to give Quebec a veto over constitutional amendments. After the referendum was won, however, it was clear that none of the other provinces would approve of a unilateral veto given to Quebec and so what was eventually developed was a so-called 'Five-Region Veto.' This was essentially a form of weighted voting in which the federal cabinet could not introduce a constitutional resolution unless it already had the approval of Ontario, Quebec, British Columbia, two of the three prairie provinces consisting of at least 50% of the prairie population, and two of the four atlantic provinces consisting of at least 50% of the population of that region. Under the old '7/50' system however, all that was needed to introduce a constitutional amendment was approval by any coalition of seven provinces consisting of 50% of the population or more.[7] When we use the BPI to

quo, the out come of a vote, and the variety of intermediate states preferred by the various actors, are plotted as points in a finite space, where the distance between an actor's preferred state and the outcome of a vote is used as a measure of their preference. As such, a more accurate representation of player interactions is developed. Yet this method is beyond the scope of this paper as it does not directly concern the BPI.

[7] There are a lot of interesting side-effects to this change in voting system, such as making Alberta a de facto veto holder as it has over 50% of the prairie population and is therefore always a critical member of a winning coalition, and conversely making Prince Edward Island a 'dummy' player with no influence

calculate and compare the distribution of relative voting power under both these systems we see that Alberta is actually the big winner with an increase of 70% in voting power, followed by British Columbia at 57% and Quebec at 43%. The big losers on the other hand, are Prince Edward Island which lost 100% of its voting power, followed by Manitoba and Saskatchewan who both fell to 57% of their original values.[8] In general then, according to the BPI in this case some provinces benefited from the change while others did not. Yet this view omits some very important data. Since the BPI is calculated in part by summing up the number of times a player can block a winning coalition by defecting from it, it provides a one-sided view of voting power. The winners under the new system are winners only in a negative relative sense in that they can block more coalitions through the use of their veto, but the BPI tells us nothing about any absolute change in power, such as the ability to bring about change or introduce a new amendment to the constitution. To calculate each province's positive power to achieve an amendment and determine the absolute change in power, one must multiply the ratio of all winning coalitions to the total number of possible coalitions by each province's BPI. The result in this case is that by switching from the old 7/50 system to the Five-Region Veto, the number of winning coalitions dropped from about 16% to 3%, and therefore every province is a loser in the sense that their positive power to bring about change has dropped significantly.[9] In other words, while the BPI indicates that Alberta may have benefited from an increase in its relative ability to block resolutions, when we factor in the new ratio of winning coalitions to total coalitions we see that it has actually decreased

---

whatsoever as its population is so small it will never be a critical member in a winning coalition of atlantic provinces.

[8] Flanagan, p.101

[9] Flanagan, p.102

in its ability to obtain future changes to the constitution. As a result, the true consequence of the change in voting system in the Canadian government was not to increase the overall power of neither Alberta nor Quebec but rather to simply preserve the status quo, an effect not predicted by the BPI as it has no capacity to consider absolute power.

In addition to these aforementioned limitations of the BPI, it should also be noted that there are a few drawbacks related to the very nature of an index which limit them as a method of evaluating power within a group; it is difficult for any index in general to encompass either the complexity or the variability of a an institutional voting system. A power index is often too simple to take in to account the complex nature of real-world voting systems, such as those that take place in the intricate processes of the European Union, in which a complex interplay between the Commission, the Council of Ministers, and the European Parliament is needed to pass a resolution. Additionally, if we accept power as defined as a player's ability to influence others, then it becomes a function of that player's votes, the number of votes held by the other players, the decision rule needed to pass a resolution and any changes that occur to any of these values. Consequently, the relative power in the group will fluctuate significantly given changes in characteristics like the number of players in the group or a change in whatever factor is used to assign votes to an actor, such as a change population. An index, however, is static and can only give a relative assessment of a given set of these values rather than a fluid change over time.

The BPI therefore can in general only be used as a rough measure of the distribution of power in a group, and only according to the current *a priori* decision making rule. By assuming that all coalitions are of equal value the BPI disregards the impact of P-power on a group, an impact that can be quite significant as illustrated by the JPI, and by assuming that all coalitions are equally likely to occur, the BPI additionally disregards the impact of player preferences on the distribution of power within the group, thereby eliminating a significant aspect of real-world power dynamics – the interaction of the players. Furthermore, in the event that the decision making rule is changed, this index can only address any subsequent change in relative power within the group, but can not reveal any change in absolute power, as illustrated by the freezing of the status quo in the Canadian constitution following a change in the voting system there. Consequently, the Banzhaf Power Index, as a tool for assessing voter power in a weighted voting system, suffers from several limitations arising from its innate assumptions and inabilities, and can not provide a full and accurate measurement of the distribution of power in a system, but to do so must instead be employed in conjunction with other tools of power measurement.

References

Braham, Matthew & Holler, Manfred J. (2005). The Impossibility of a Preference-Based Power Index. *Journal of Theoretical Politics*, 17(1), 137-157.

Felsenthal, Dan S. & Machover, Moshe (2005). Voting Power Measurement: A Story of Misreinvention. *Social Choice and Welfare*, 25(2-3), 485-506.

Felsenthal, Dan S. & Machover, Moshe (2001). Myths and Meanings of Voting Power: Comments on a Symposium. *Journal of Theoretical Politics*, 13(1), 81-97.

Flanagan, Thomas (1998). Game Theory and Canadian Politics. Toronto, ON: University of Toronto Press.

Strand, Jonathan R. (2005). Regionalizing Multilateralism: Estimating the Power of Potential Regional Voting Blocs in the IMF. *International Interactions*, 31(1), 15-54.