King Abdullah Petroleum Studies and Research Center

Multidimensional Bargaining Using KTAB

Ben Wise, Leo Lester and Brian Efird

February 2016 KS-1524-DP018B

About KAPSARC

The King Abdullah Petroleum Studies and Research Center (KAPSARC) is an independent, non-profit research institution dedicated to researching energy economics, policy, technology, and the environment across all types of energy. KAPSARC's mandate is to advance the understanding of energy challenges and opportunities facing the world today and tomorrow, through unbiased, independent, and high - caliber research for the benefit of society. KAPSARC is located in Riyadh, Saudi Arabia

Legal Notice

How to use this Document

This paper (*Multidimensional KTAB*) is a technical discussion paper designed as a follow on to *An Introduction to the KAPSARC Toolkit for Behavioral Analysis using one-dimensional spatial models (Unidimensional KTAB)*. It extends the original framework of KTAB to explain the analysis described in the recently released KAPSARC discussion paper *Reforming the Role of State-Owned Enterprise in China's Energy Sector: An Analysis of Collective Decision-Making Processes Using the KAPSARC Toolkit for Behavioral Analysis (Chinese SOE Reform).*

In the first paper, we introduced the basic concepts of the KAPSARC Toolkit for Behavioral Analysis (KTAB), what we mean by Collective Decision-Making Processes (CDMPs) and the basic terms surrounding simple models implemented in KTAB, such as the Spatial Model of Politics (SMP), including actors, position, influence and salience.

In *Unidimensional KTAB*, our aim was to introduce these basic concepts by focusing on what we called simple dynamic CDMPs. This type of CDMP describes a situation in which actors could generate a series of proposals and counterproposals in an attempt to win enough support from the rest of the group until no actor can improve their position, which determines a final winning proposal.

The actors can change their positions, but not their utility functions. The first paper also introduced the concept of complex dynamic CDMPs (*Unidimensional KTAB*: section 4.2, page 19). These were considered to be an extension to the simple dynamic case, in which actors can induce other actors to change both their position and utilities. These complex dynamic CDMPs are the topic of this subsequent paper. In addition, we expand our explanation from one-dimensional problems to multidimensional questions.

This paper aims to provide applied users with an understanding of the logic and theory underpinning the models and, more specifically, the underlying mechanics behind the case study presented in *SOE Reform*. It also explains to power users how to build and construct models of multidimensional CDMPs using the SMP in KTAB.

For applied users who wish to run their own examples through KTAB using the same model, the software libraries are available on the KTAB website at http://ktab.kapsarc.org. Users can run their own models with **smpc**. It can read a user-specified CSV data file for a scenario; for testing purposes it can also generate random test scenarios. More information on how to run the **smpc** is available in the online documentation. Source code documentation and other supporting materials are available through the KTAB website.

Table of Contents

1. Expanding the Spatial Model of Politics

In *Unidimensional KTAB*, we introduced:

- a one-dimensional CDMP model known as the Spatial Model of Politics (SMP), and
- some basic terminology necessary to describe the CDMP model and collect the necessary data.

In this section, we both reiterate the main points from that paper and add an extra layer of sophistication necessary to handle the subtleties of a multidimensional CDMP. We begin by restating what an actor is and explaining the expansion of the SMP from one dimension to many.

1.1 Actors

Actors are all the stakeholders that contribute to the CDMP in some way. Actors are generally groups of individuals or entities with common interests. When actors are labelled with the names of individual leaders, they are taken as representative of the faction which that individual leads. Aggregates can be formal (such as a corporation) or informal (such as ‘subsistence farmers’ or ‘aerospace manufacturers’).

We will generally use the symbol N to denote the number of actors and refer to each individual actor by an index such as i, j, or k. The more exhaustive the list, the better the position and the power landscape will be mapped, but there is of course a trade-off between an exhaustive list and the available knowledge.

1.2 Moving from the One dimensional SMP to a Multidimensional SMP

In *Unidimensional KTAB* we introduce various concepts, including the Practical Spectrum of Plausible Positions (PSPP). This was a one-dimensional line that enabled us to convert qualitative descriptions of policies into a numeric position. Figure 1.1 recreates the example spectrum of positions used in *Unidimensional KTAB* (section 3.3, page 14).

On this one-dimensional spectrum, actors advocating a complete free market would be given a position of 0; actors calling for complete government control, 100. An actor calling for liberalized markets in consumer goods, but with government control over education and healthcare, might be given a position of 25 or 30.

Not all questions and CDMPs can be so neatly reduced to a single dimension. A broader policy discussion over state control could be imagined. State control of markets could be just one aspect; another might be state control over the press and freedom of speech. It is possible that a single actor could advocate liberalized markets but high levels of censorship, or vice versa. The point is that actors can view different dimensions of the overall policy differently. This is the basis of the expansion towards the multidimensional SMP.

1.2 Moving from the One dimensional SMP to a Multidimensional SMP

A two-dimensional SMP might therefore be visualized as in Figure 1.2, with the different policy dimensions set out on different axes, where each axis is a well-defined PSPP.

Of course, this might be extended to three dimensions through the addition of a third explicit policy option along a third axis, labeled as a PSPP (Figure 1.3), and from there to the generic multidimensional case with k dimensions.

As the dimensional space of the SMP becomes more complex, it is necessary to define an actor's position with increased clarity and specificity.

In *Unidimensional KTAB*, we dealt with only a single dimension (the PSPP), and each actor's position was a single number between 0 and 1. However, a multidimensional policy debate is one in which there are multiple, separable aspects of the problem, so it takes multiple co-ordinates to describe the position of an actor.

Positions are points in multiple dimensions. Each aspect of the problem is assigned its own dimension, and the actor's preferred outcome for that aspect is its co-ordinate along that dimension. For convenience, each co-ordinate is normalized to the scale. In a policy debate with three issues, an actor's position would be described by a triple of three co-ordinates, e.g. (0.23, 0.98, 0.67).

A two-dimensional example is given in Figure 1.4. There are five actors, A through E, whose positions are represented as points inside a two-dimensional square. Each actor's position is given by two co-ordinates, one for the *X*-axis and one for the *Y*-axis. Note that actors can be close on one dimension while being far apart on another, as with actors A and B. It is important to notice that the

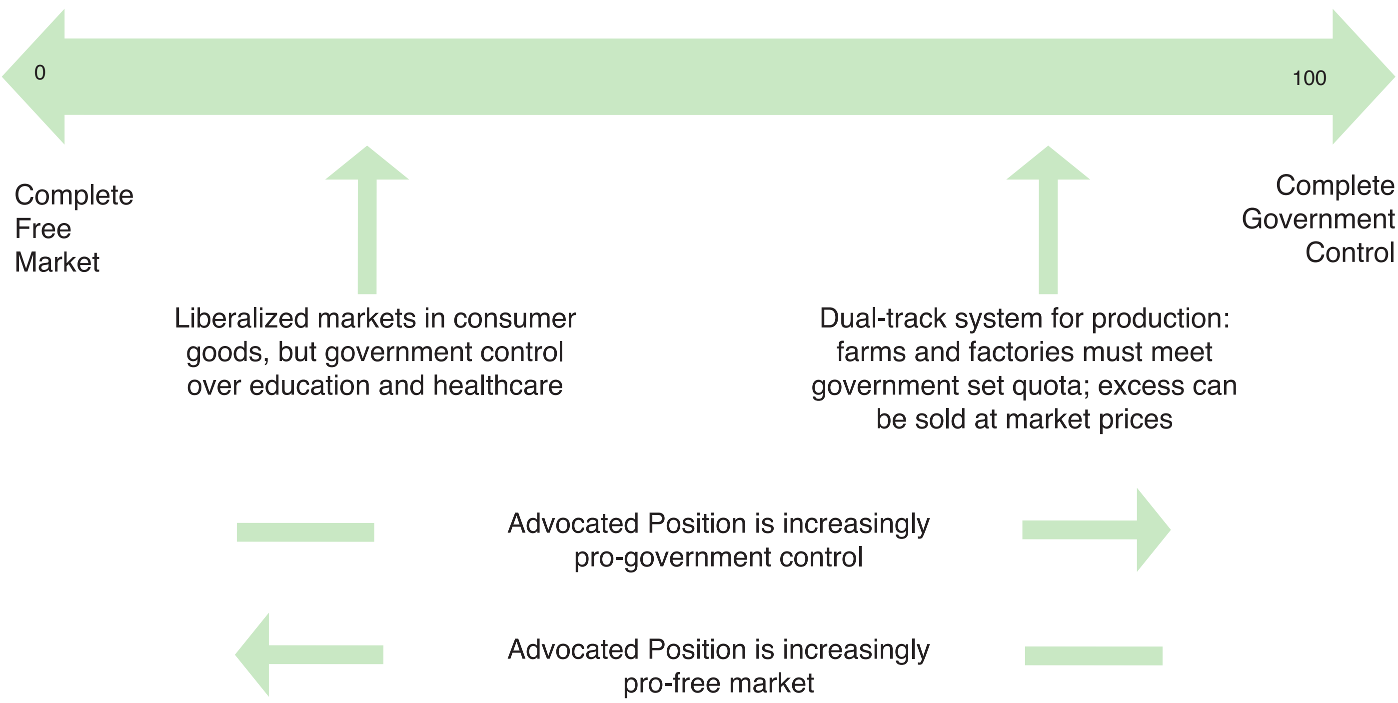

Figure 1.1 - A simple one-dimensional spectrum of positions, using economic policy as an example (recreated from Figure 3.1 in Unidimensional KTAB*: section 3.3, page 14). The spectrum displays attitudes with regard to a single question of issue with movement in either direction signifying incrementally more extreme views. Source: KAPSARC*

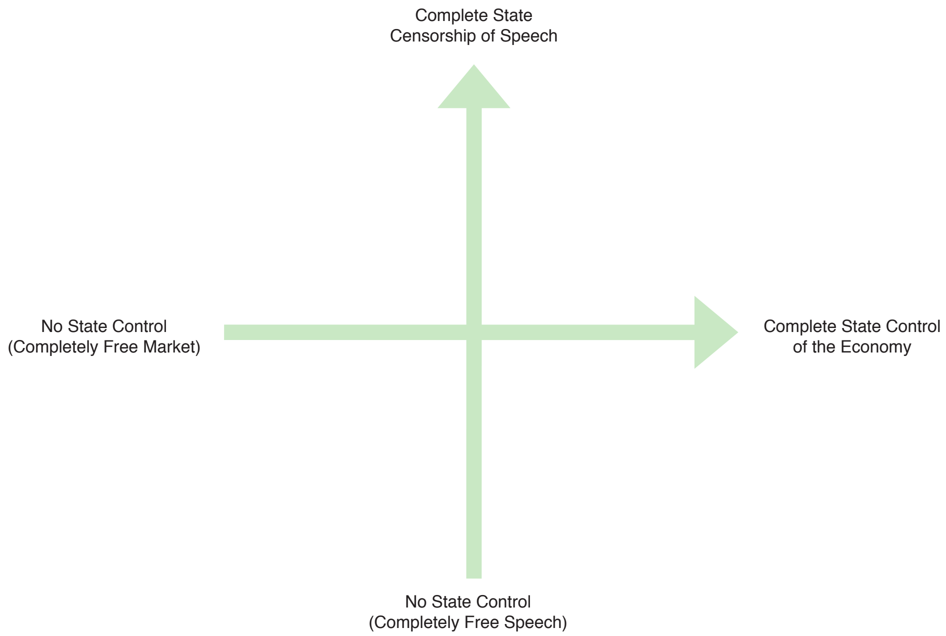

Figure 1.2 - Extending the Policy Issue to Two Dimensions. The initial PSPP describing economic policy remains as before. A second dimension has now been added as a Y-axis looking at press censorship and freedom of speech. Again, more extreme positions are found towards the ends of the axes. Actor positions are no longer points on a line, but points in two-dimensional space. Source: KAPSARC

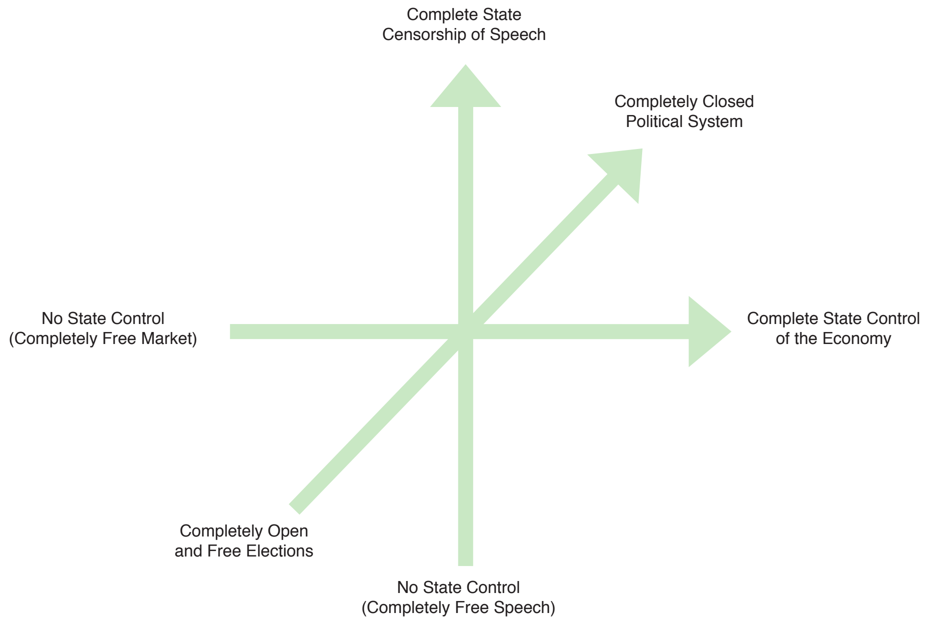

Figure 1.3 - Extending the Policy Issue to Three Dimensions. A third axis has been added to describe a third element. Positions towards the overall policy are now described as points in three dimensional space depending on attitudes to the three separated policy dimensions. Source: KAPSARC

actor's position is a point in two-dimensional space; where an actor falls along a single dimension is simply a co-ordinate. The position is the list of all the co-ordinates along all the PSPPs.

Sometimes, a policy debate may be conducted in terms of many specific issues, but the positions on specific issues are highly correlated with each other-because they are all just aspects of one underlying conceptual issue. For example, a great many domestic economic debates are about different aspects of income redistribution. While there are many nuances to the debate, knowing that an actor favors extreme and thorough redistribution of wealth from rich to poor is sufficient information to infer its position on many particular policies for income redistribution. Similarly, knowing that an actor is totally opposed to any income redistribution is sufficient information to infer its position on many particular policies for income redistribution. In statistical terms, this means that there is one 'principal component' (pro- or anti-redistribution) that explains most of the variation in the data. Principal component analysis is often used in public opinion survey data to try to summarize the underlying issues driving policy debates; for our purposes it is a useful concept to reduce the dimensionality of the problem.

This is illustrated in Figure 1.5. Although the blue data points were originally obtained in terms of *X* and *Y* dimensions, it be might better to analyze the data in terms of the green principal component. This is frequently the case when two or more sub-issues are not really separable, because they (*X* and *Y*) are different components of a plan to implement a policy, or different aspects of the consequences expected from a policy. In this case, the analysis would be best conducted by abstracting away the details and focusing on the underlying issue.

In short, we may find that what was originally thought to be a two-dimensional problem can be more succinctly modeled as a one-dimensional CDMP.

We will generally use K to denote the number of dimensions in a policy debate. The set of all possible positions is $\Omega = [0,1]^K$. A particular position will be referenced as another Greek letter like θ, α, or β. The position of the i-th actor would be θ_i and the k-th co-ordinate would be θ_{ik}. Thus, the complete position is represented as an ordered list of all the co-ordinates along all the PSPPs:

$$\theta_i = (\theta_{i1}, \theta_{i2}, \theta_{i3}, \ldots, \theta_{iK}) \quad (1)$$

1.3 The Main Terms in the SMP

As we move from the one-dimensional SMP to our current multidimensional case, it is necessary also to redefine the main terms.

Influence measures how easily the actor can shape the outcome of the CDMP if fully motivated. This is not a measure of how likely the actor's preferred position is to win, nor is it a measure of the actor's motivation to win. It is the maximum capability to exert influence, clout, or political power which an actor could apply to the policy debate. It could be a consequence of its political connections, official or unofficial position, ties to critical stakeholders, economic wealth or other factors.

In *Unidimensional KTAB* (section 5.4, page 27) we denoted influence with the letter j. As we develop the sophistication with which an actor's influence can be modulated, we will generally use c_i to denote the i-th actor's maximum capability to exert influence; it is always positive. The basic purpose of the capability measure is to assess whether one

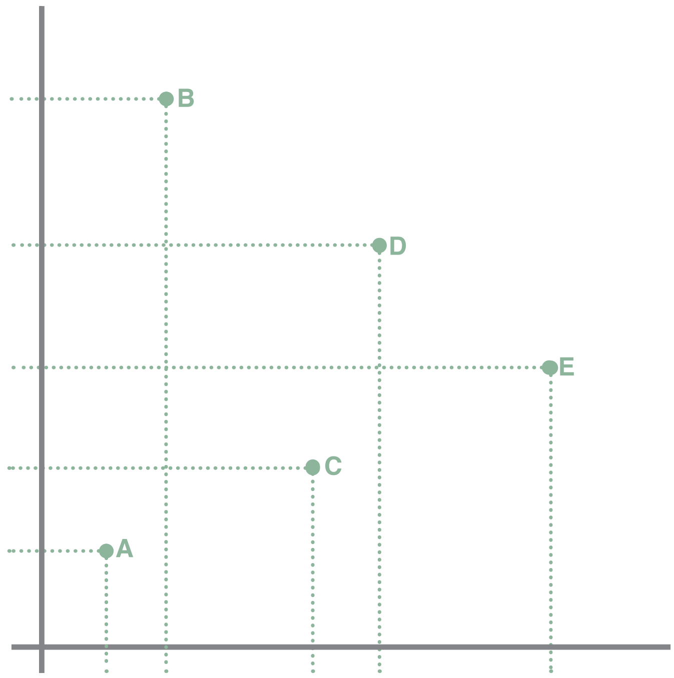

Figure 1.4 - Positions in Two Dimensions. Each actor has an attitude towards each dimension of the policy question (axis co-ordinates), but its overall position towards the whole policy question is denoted by a point in multidimensional space.
Source: KAPSARC

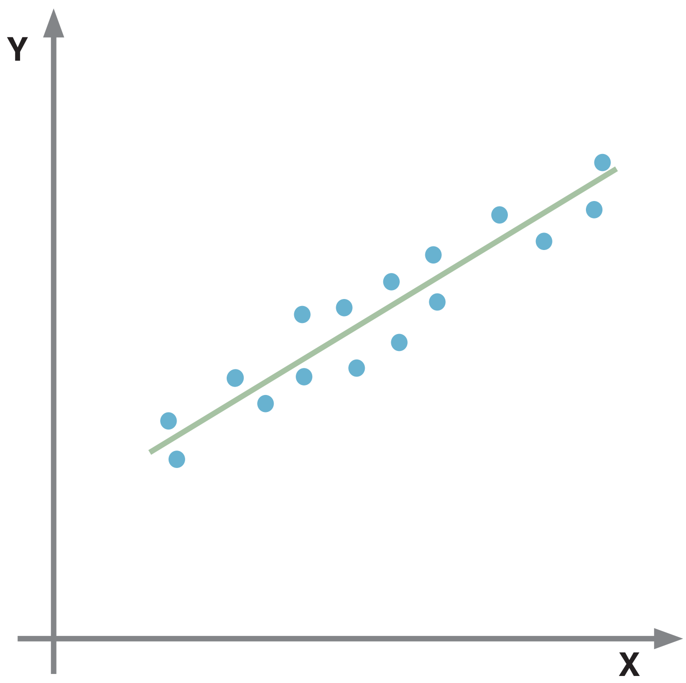

Figure 1.5 - Correlated Co-ordinates. Highly correlated dimensions may reveal an underlying principal component, in which case it is better to reduce the number of dimensions to focus on the principal component (in this case from two to one).
Source: KAPSARC

coalition of actors is stronger than another coalition in the policy debate. As only the ratios of capability matter in comparing coalition strength, it is often convenient to normalize the values so that the strongest actor has maximum influence, 100.

In *Unidimensional KTAB* we introduced the term Attractiveness to describe how actors view the desirability of the expected consequences of various positions. In the one-dimensional case that we described, there was no distinction made between the position taken and its expected consequences. This means that there is no difference between the way the actor views the position and the way the actor views the expected consequences of that position. They are one and the same thing.

In the multidimensional case, this agreement on how positions translate into consequences is lost. Each actor could have a somewhat different mental model of the world. This means that positions can no longer be thought of as synonymous with expected consequences.

Utility is the desirability of a position to an actor, typically normalized to a [0,1] scale.

Therefore, to repeat:

> Attractiveness is a numeric function of the expected consequences of an outcome;

> Each actor has a world view or model (WM) which enables them to convert a position into expected consequences according to their own perspective; and

> Utility is a numeric function of the position. The utility of the policy (position) is the attractiveness of the expected consequences of that policy.

The utility applies to the position as a whole, taking into account all the co-ordinates along each dimension. While the utility measure always takes into account the interactions between dimensions, it is often useful to analyze how changing just one dimension of a particular multidimensional position would change an actor's utility.

As Figure 1.6 shows, utility need not change equally in each direction: it is a function of a multidimensional position. Nonetheless, we can examine the utility of each dimension separately by sliding the position *A* (with utility U_A) along a single dimension, to give a new position of either *A*' (with utility $U_{A'}$) or *A*" (with utility $U_{A''}$).

Finally, **salience** is how much importance an actor attaches to each dimension of a policy. In one dimension, an actor's salience was taken to be the overall importance to that actor of the entire issue, as there was only a single issue. In simple terms, it described the likelihood of the actor to focus on the issue rather than being distracted by competing concerns outside the modeled CDMP (see Figure 1.7). Salience has previously been described in a scale [0,100]; as with utility it can also be normalized to [0,1].

In multiple dimensions, it will often happen that different actors will place differing emphases on different aspects of the problem. Different dimensions will therefore have different salience scores describing their relative importance to the actor; see Figure 1.8.

Consequently, the i-th actor has an entire vector of salience values, where s_{ik} is the salience to the i-th actor of the k-th policy dimension. The overall salience for actor i is the sum s_i over all dimensions of the salience to that actor of that dimension:

$$s_i = \sum_{k=1}^{K} s_{ik} \tag{2}$$

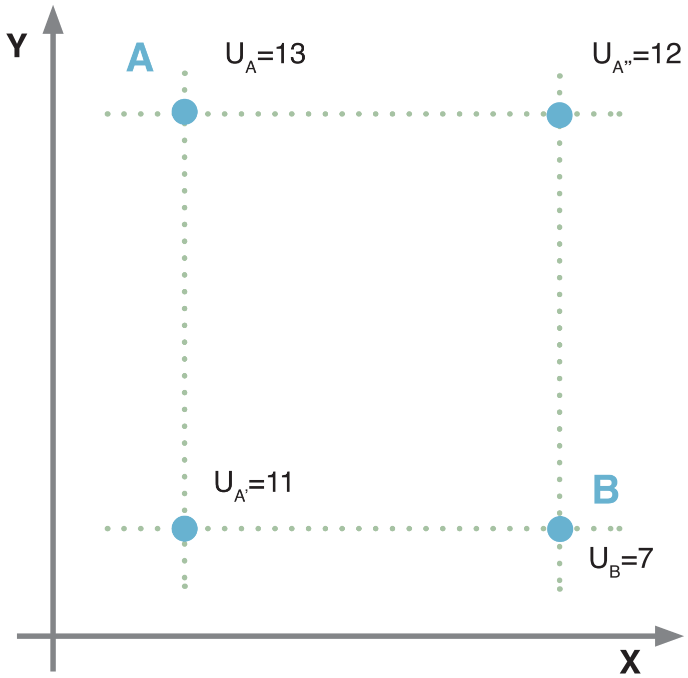

Figure 1.6 - Analyzing the utility of a point in two-dimensional space. The utility of a position is a function of the precise trade-offs described by the point in multidimensional space. Moving a point along a single dimension (e.g. U_A to U_A') can reveal the importance of that dimension, but it will not fully explain the relative and interacting importance of the two dimensions in explaining the utility of a different point in space (U_A vs U_B) Source: KAPSARC

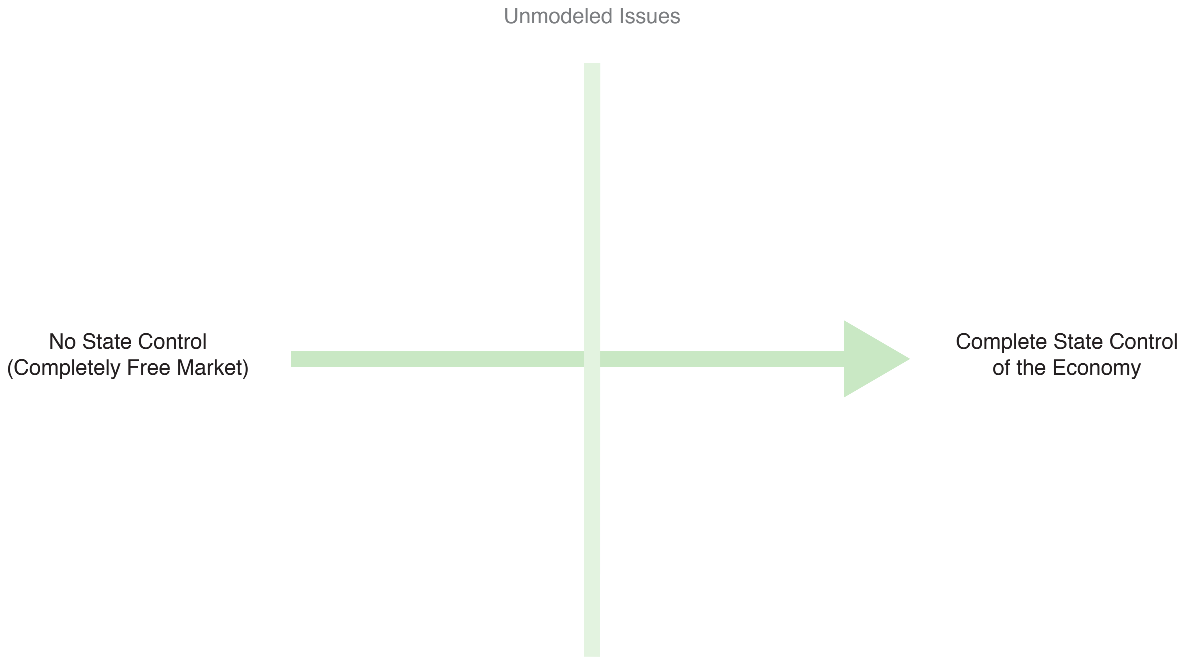

Figure 1.7 - In one dimension, salience describes how important the modeled dimension is compared with unmodeled issues. Source: KAPSARC

This equation does not alter the meaning of overall salience; it simply describes its components. Total salience will sum to less than or equal to 1; the difference from 1 reveals the importance of unmodeled aspects.

Another way of thinking about salience along one dimension is to view it in terms of the stakes for changing only that dimension. The higher the stakes of a dimension to an actor, the higher the salience of that dimension. The situation can be illustrated in Figure 1.9.

The blue dot marks the position of an actor. As this is its own favored position, this has utility 1.0 to it (the actor). Positions further away have lesser utility; the green contour lines show how the utility declines with distance. Because the contour lines are not circles but ellipses, changes in the *X* co-ordinate have more impact on its utility than do similar changes in the *Y* co-ordinate. In other words, it will perceive higher stakes from changes in *X* than from similar changes in *Y*. In other words, *X* has higher salience than does *Y* for this actor.

It is important to understand that the influence actually exerted by actors in a choice between two alternatives depends directly on the stakes, and only indirectly on the distance. Referring back to Figure 1.9, the light blue dot at the top has utility 0.7 to the actor while the yellow dot to the right has utility 0.4 – even though both are the same distance from the preferred point (the blue dot). Because *X* is more salient to this actor than is *Y*, the same distance in the *X* direction is of more consequence than the same distance in the *Y* direction. The stakes in a

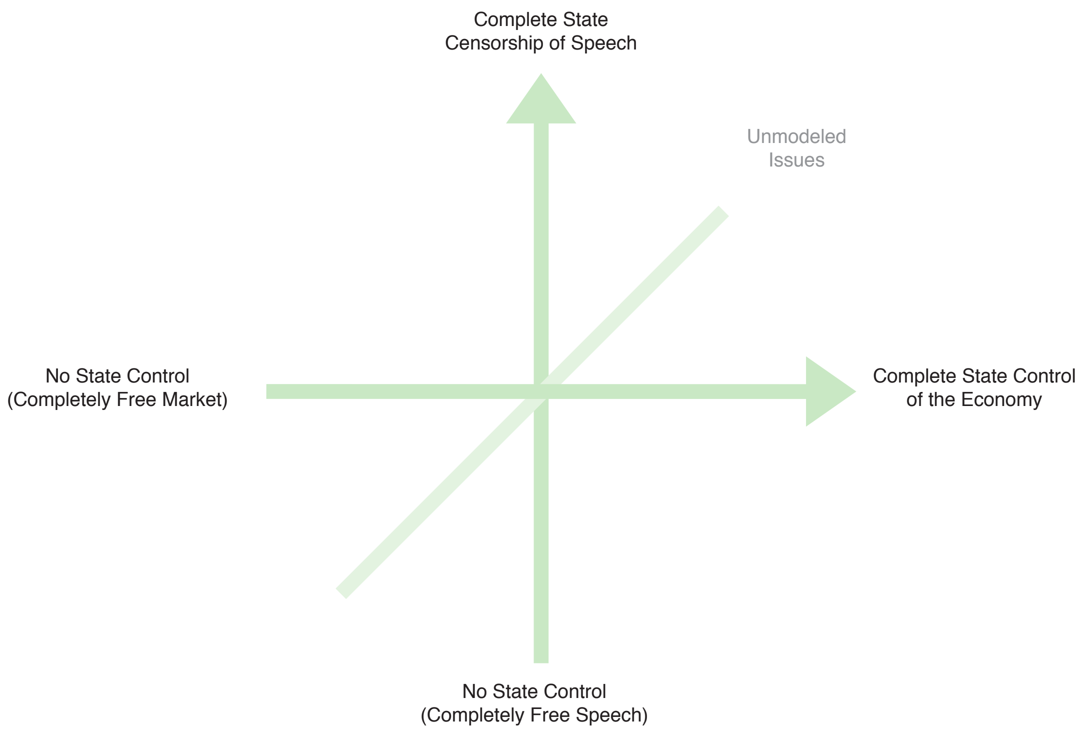

Figure 1.8 - In multiple dimensions, overall salience describes how important the modeled dimensions are compared with unmodeled issues. Each dimension has its own individual salience. Source: KAPSARC

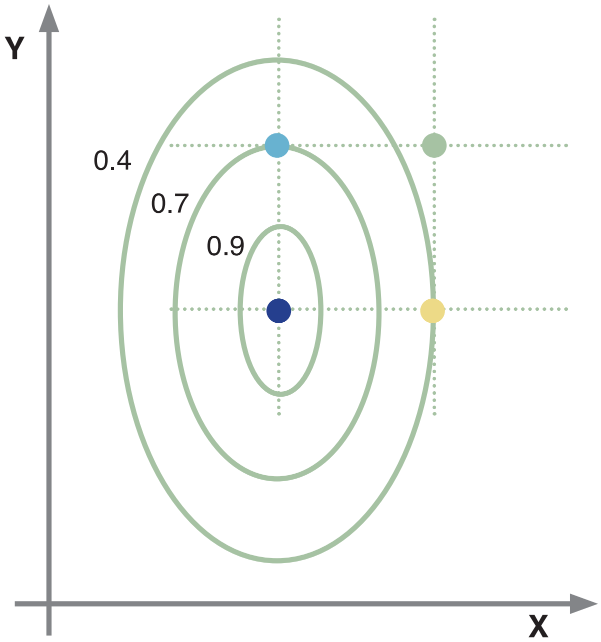

Figure 1.9 - Utility Contours. The light blue and yellow dots are equidistant from the blue dot but the different salience scores assigned to each axis means that the move from blue to yellow represents a larger loss of utility than the move from blue to light blue. Source: KAPSARC

choice between blue and light blue would be 0.3 = 1 - 0.7, while the stakes in a choice between blue and yellow would be 0.6 = 1 - 0.4. The actor would be much more motivated to act in the second case and could be expected to exert roughly twice the effort.

Notice that the green contours are evenly spaced along the *Y*-axis, but the utility continues to decline. The decline from the blue dot to the first ellipse is just 0.1; the decline from the first to second ellipse is 0.2; the decline from the second to the third ellipse is 0.3. The shape is actually an inverted parabola (upside down letter U), centered on the blue dot and stretched in the *Y* direction.

The **stake** in a choice between two positions is the difference between what an actor sees as the utility between those positions. The stake is only defined in terms of the comparison between two different options. Suppose an actor is faced with four options, where A and B are both very desirable, while C and D are both very undesirable. It would have very little at stake in an A:B choice, because both are very good options, so it would probably not be greatly concerned whichever occurred. However, it would have a great deal at stake in A:C or B:C, because C is much worse than either A or B. Finally, it would have little at stake in C:D, because both options are equally bad.

In Figure 1.10 we show this graphically: d_1 and d_2 have similar levels of utility for the actor, regardless of their apparent distance on the PSPP. The stakes are therefore low and the actor will likely not care which of d_1 and d_2 is the outcome. In contrast, d_2 and d_3, although seemingly close on the PSPP, have very different levels of utility. In the choice between d_2 and d_3, the actor is likely to fight hard to avoid d_3 as the stakes are now much higher. This reinforces the distinction between utility and proximity.

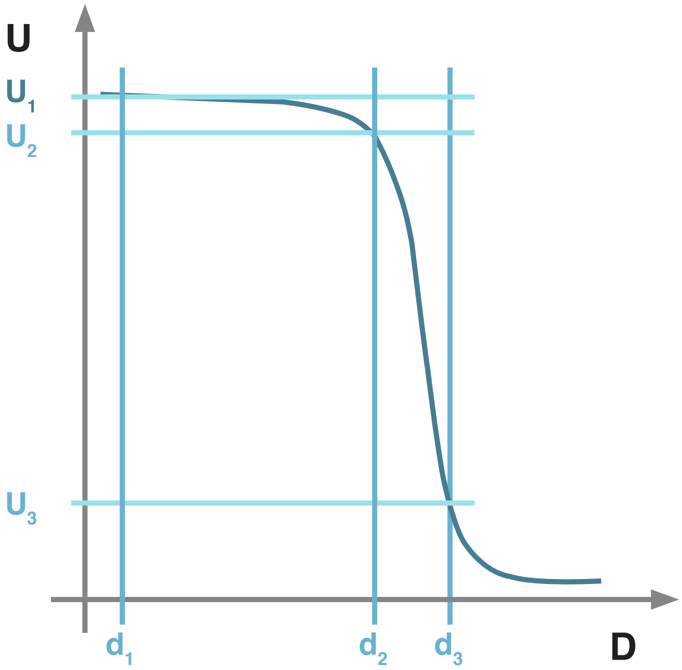

Figure 1.10 - Changes in utility need not be linearly related to changes in distance or proximity. Stakes describe this difference in utility between two positions. The X-axis (D) denotes distance between positions (in one-dimensional terms, think of this as the distance between points on a PSPP). The Y-axis denotes the utility of these positions. Positions 1 and 2 have similar utility despite their distance: low stakes. Positions 2 and 3 have very different utilities despite their closeness: high stakes.
Source: KAPSARC

What an actor has at stake strongly affects how much influence it will exert to shape the outcome of a policy. When a 'life or death' issue arises, with the highest possible stakes, great effort can be expected. When a trivial issue arises, with insignificant stakes, little effort can be expected. Of course, what is critically important to one actor may be unimportant to another – and vice versa.

> Stakes and salience are very closely linked but are, importantly, conceptually distinct.

> The stakes refer to the difference in utility between two positions in multidimensional space.

> Unlike stakes, salience does not refer to how the actor will choose between pairwise alternatives but rather:

- The overall salience, s_i, refers to the degree to which an actor cares about the overall policy discussion being modeled.
- Individual salience, s_{ik}, refers to the weight the actor places on a particular individual dimension of the multidimensional policy question.

1.4 How the Terms are Related

Influence, utility, stakes and salience are all tightly interconnected, with the latter three all modulating, in one way or another, the degree of influence an actor actually exerts at any particular moment during the CDMP. One way of formalizing these relationships is the following Domain Specific Utility Model (DSUM).

The first part of the DSUM is to measure the policy relevant difference between two positions in terms of an actor's salience. That is, two positions, α and β may have different co-ordinates on the k-th dimension, separated by a certain distance, $\alpha_k - \beta_k$, but if that dimension is not important to the actor, then it sees no significant difference between the policies. Another actor, who places higher salience on that dimension, might see the same distance as being highly significant.

$$d_i^2(\alpha, \beta) = \frac{\sum_{k=1}^{K} \left(s_{ik}(\alpha_k - \beta_k)\right)^2}{\sum_{k=1}^{K} (s_{ik})^2} \qquad (3)$$

There are several important features of equation 3. The denominator cannot be zero: the actor must regard some aspects of the policy as having at least some significance. If not, that actor probably should be removed from the model. Second, the difference is essentially an average, weighted by salience. Because all the positions are in the K-dimensional space $[0,1]^K$, the separation between any two co-ordinates is between 0 and 1, so the overall policy relevant difference measure is between 0 and 1.

Notice that if the actor does not care about distances along the k-th dimension, formally $s_{ik} = 0$, then the separation, $\alpha_k - \beta_k$ has no effect on the difference d_i. When we say that an actor "has no position" on the k-th dimension of a policy problem, we really mean that it has no position that it is willing to promote or defend, because it has no preferences along that dimension, i.e. that $s_{ik} = 0$. One consequence of $s_{ik} = 0$ is that when it is time to bargain over positions, actor i will make no attempt to preserve θ_{ik} and will acquiesce to whatever any other actor proposes along that dimension, because it simply does not care.

This situation is shown in Figure 1.11, where the dimension has zero salience. It is similar to figure 1.9, except that the ellipses have been stretched until their sides appear as two parallel lines. The inner pair correspond to the positions with utility 0.9, the middle pair to positions with utility 0.7, and the outer pair to 0.4. Now, because *Y* has zero salience and utility does not change at all along that axis, the orange position has the same utility, 1.0, as does the actor's position at the blue dot.

This gives an interpretation of the overall salience: it is the salience of the policy problem being modeled, compared with all the other (unmodeled) factors affecting that actor. Thus, the overall salience of the actor measures its propensity to attend to this policy issue, as was explained in *Unidimensional KTAB* (section 3.5, page 17). Conversely, 1-s_i is a rough indicator of the actor's propensity to acquiesce in proposed changes to its overall policy position, because it is attending to other factors. In other words, 1–s_i is the salience of all the external factors which are not directly represented.

In the one-dimensional case $K = 1$, $d(\alpha, \beta) = |\alpha - \beta|$ regardless of s_i. However, in the multidimensional case, actors will typically disagree as to the policy relevant difference, because they attach different salience to different aspects of a policy. In symbols, $d_i(\alpha, \beta) \neq d_j(\alpha, \beta)$.

The second part of the DSUM is to treat the utility to an actor of a position as declining with the difference between the actor's position and that position. Again policy relevant difference is not the same as simple distance and depends on the particular actor being considered. Each actor has a curvature parameter *R*, the meaning of which will be explained shortly.

In multiple dimensions, it is likely that actors will not agree on the order in which policies that are more or less similar to their own should fall, because each has their own idea of the policy relevant difference. This effect is illustrated in Figure 1.12. As before,

1.4 How the Terms are Related

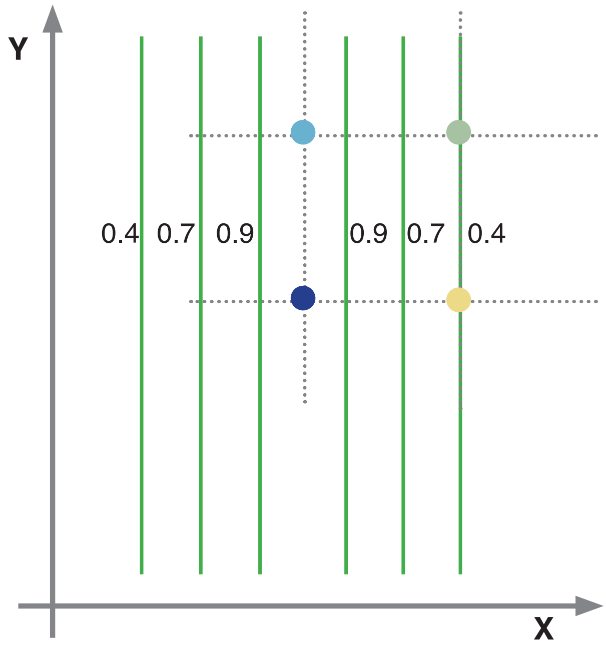

Figure 1.11 - Zero Salience Contours. The actor has a salience score of zero for the Y-axis. As a result of zero salience the stakes between positions that differ only in their Y-axis co-ordinate are also zero. The actor is agnostic between the blue and light blue dot. Source: KAPSARC

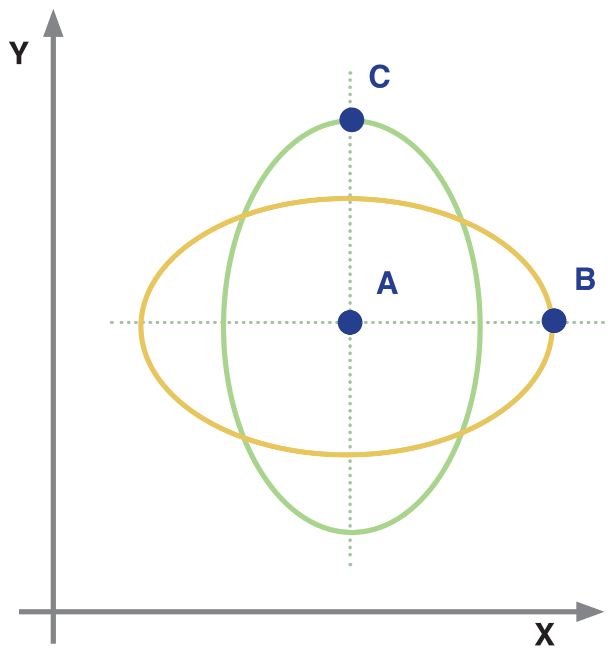

Figure 1.12 - Different Orders. Different actors have different views of the world by virtue of their difference salience scores. Green would rank positions [A,C,B]; Yellow would rank them [A,B,C]. Source: KAPSARC

Influence, utility, stakes and salience are all tightly interconnected, with the latter three all modulating, in one way or another, the degree of influence an actor actually exerts at any particular moment during the CDMP.

the green ellipse is the utility contour for an actor whose own position is the point *A* and which has higher individual salience for the *X* dimension for *Y*. By contrast, the yellow ellipse is the utility contour of an actor the position of which is also at *A* but which has higher individual salience for the *Y* dimension than for *X*. The ellipse has the same shape except rotated: both actors have the same overall salience, except their individual saliences are swapped between *X* and *Y*.

If the green actor was asked to order the three blue dots by how similar they were to his position at *A*, then his order would be [*A,C,B*]: the green ellipse would have to shrink down to reach *A* and would have to stretch out to reach *B*. However, if the yellow actor was asked to order the dots by how similar they were to his position *A*, his order would be [*A,B,C*].

Thus it is perfectly possible in multiple dimensions that two actors could have the same position, the same overall salience and the same capability to exert influence, and still have dramatically different preferences between options because their individual saliences were different. This is not possible in one dimension, because overall salience and individual salience are the same in one dimension.

As suggested in Figure 1.12, the utility function of an actor is typically nonlinear; the degree of nonlinearity is given by the curvature parameter *R*. The utility function is a quadratic function of difference:

$$u_i(\alpha) = \big(1 - d_i(\alpha, \theta_i)\big) \times \big(1 + R_i d_i(\alpha, \theta_i)\big) \qquad (4)$$

Without subscripts, this can be expressed more succinctly as $u = (1 - d)\,(1 + Rd)$. This is the DSUM that gives elliptical utility contours in Figure 1.9.

2. Introducing a Generalized Approach for Modeling CDMPs

In *Unidimensional KTAB* we introduced some basic concepts to explain the modeling of CDMPs, but in fact we were discussing:

- one-dimensional questions; and
- deterministic answers.

Black's (1948) celebrated Median Voter Theorem (MVT) says that the winning outcome of a CDMP is that which has 50% or more of the vote to its right and 50% or more of the vote to its left. Left and right here refer back to the horizontal one-dimensional PSPP upon which the analysis is based.

The MVT suffered from various limitations:

- it allowed for only binary voting: a simple yes or no, with the actor's full voting weight;
- it only worked for one-dimensional questions; and
- it assumed a single-peaked utility curve for each actor.

For a fuller discussion, see *Unidimensional KTAB*, section 4, page 18.

The Central Position Theorem (CPT) was then introduced as an alternative to overcome the limitations of these restrictive assumptions:

- proportional voting was allowed instead of binary voting;
- multidimensional and even nondimensional questions can be modeled; and
- actor utility curves can take on a variety of forms.

In all cases, though, the theory as presented in *Unidimensional KTAB* led to deterministic outcomes. The winner of any CDMP was declared the absolute winner regardless of the margin of his victory. In practice this is not always the case. We can imagine some scenarios where the vote is carried by a majority of one, but we can also imagine scenarios where having a marginal increase in voting weight over an opponent does not equate to certain victory. An example could be a lobbying campaign in which having more money can certainly improve the chances of winning, but need not seal victory. Poor use of the money could enable the weaker, more poorly funded, lobby group to prevail. In other words, we need to develop a more nuanced view of winning, moving away from a deterministic outcome to a probabilistic one. In this section we will explore the basic theory and algebra necessary to build a probabilistic CDMP model. We call this the Probabilistic Condorcet Election (PCE).

2.1 The Probabilistic Condorcet Election

Faced with a choice between two options, the amount of influence which the actor i will exert to promote the first option over the second is termed the 'generalized vote' of actor i.

It is written as:

$$v_i\,(\alpha\text{: }\beta) \tag{5}$$

Positive values mean influence to promote the first option (α); negative values mean influence to promote the second option (β). The effort exerted will depend on the actor's capability to exert influence as well as the stakes they perceive.

As actors vote for one option over another, various

Testing KATB

A common question: Has KTAB been tested?

KTAB is a toolkit for building models (dynamic or static); it is not itself a model. Like all other theoretical frameworks for building models of problems, there is a confusing level of abstraction when discussing the theory per se, as distinct from explaining a particular model within the theory. The problem appears with well-known and well-established theories such as the theory of linear regression or convex optimization: there are abstract terms which can only be given concrete meaning for a particular model within the theory. For example, the fundamental equation of linear regression is the following:

$$y_i = \sum_j a_{ij} x_{ij}$$

If one were to ask what a particular value represents, the only general answer is “it depends”. Is the theory being used to model a legislative, microeconomic, pre-industrial military, astronomical, stock market speculation, macroeconomic, medical, or legal relationship – or something else entirely? What does the range of indices represent? Is linear regression “correct” in general? Again, it depends on the problem.

Some models may work in as much as they accurately describe reality. Some models may not work in that they fail to accurately describe reality. A theory can be considered useful so long as it does not preclude the construction of models that accurately describe reality. For example, linear regression is not right or wrong in general, but it can be used to produce accurate models in some circumstances. Similarly, the kind of models that KTAB can produce have certainly been used to great effect and shown to be plausible and insightful reflections of those real world CDMPs they modeled.

We introduce several abstract terms whose interpretation depends on the problem being analyzed, the most important being ‘actor’, ‘option,’ and ‘influence’. We leave these abstract terms undefined in the hope that an intuitive sense will emerge from considering how they are used in particular cases. In essence, the question of whether or not KTAB works or is correct, is more simply a question of:

- whether or not it allows one to build models that accurately reflect reality; and
- whether or not it makes it easy to construct such models.

In this instance we believe the answer is “yes”. As discussed in section 7.4 of *Unidimensional KTAB*, KTAB style models have been rigorously tested on several hundred years of data on European wars, on all US Congressional votes and several thousand policy analyses. Where the models are appropriate, they have been shown to be highly useful.

coalitions will emerge. A coalition can be seen as the set of actors voting for a particular option.

For instance, the coalition of actors supporting option α over β, that is the set of actors for which

$$v_i\ (\alpha\!:\beta) > 0 \tag{6}$$

It can be written as

$$c\ (\alpha : \beta) = \{i \mid vi\ (\alpha\!: \beta) > 0\} \tag{7}$$

The total strength of that coalition is then $s(\alpha{:}\beta)$ which can be written as the sum of the influence exerted for α over β :

$$s(\alpha\!:\beta) = \sum_{i \in c(\alpha:\beta)} v_i(\alpha\!:\beta) \tag{8}$$

The total strength of the coalition supporting option β over α is the sum of the influence exerted for β over α:

$$s(\beta\!:\alpha) = \sum_{i \in c(\beta:\alpha)} v_i(\beta\!:\alpha) \tag{9}$$

What is a Markov Process?

Markov processes are used to describe situations where what happens next depends on the current situation. Examples might include

1. The weather. Imagine that if a day had clear skies, the probability of the following day having clear skies was 70%. If the day was overcast, the probability of the following day being overcast was 55%. We could plot out a random walk through time, determining each day's weather (clear or overcast skies) according to these probabilities.

2. A running hare. Hares tend to run in circles. We can imagine a particular hare which, following a right turn will turn right again with 90% probability, but following a left turn will turn left again with 80% probability. Again, we can chart out a prediction of left and right turns for the hare based on these probabilities.

In Markov processes, the probability of jumping to a particular state (clear skies or overcast; left or right turns) depends purely on the current state (what the weather is today; what the direction of the last turn was) and ignores history (the probability of tomorrow's being overcast depends purely on whether today was overcast, regardless of the state of yesterday's skies).

In the case of a CDMP, the Markov process describes the probability with which any option (and there is no limit on the number under comparison, as long as it is finite) can emerge as the favored option, given the current set of actor strengths and preferences.

At each turn in the Markov process a probability distribution can be written to describe the probabilities of moving to each alternative state. The limiting distribution is the stable distribution of probabilities that emerges as the number of turns approaches infinity.

The probability that option α will be selected over option β can be written as

$$P[\beta \rightarrow \alpha] \tag{10}$$

and depends of the relative strength of support for it:

$$P[\beta \rightarrow \alpha] = \frac{s(\alpha : \beta)}{s(\alpha : \beta) + s(\beta : \alpha)} \tag{11}$$

At any given moment, there is a probability distribution over the options that describe the probability that any particular option will be the favored option at that moment. The favored option can change, with the likelihood controlled by the relative support of alternative options. This is a Markov process, and simple iterative algorithms to calculate the stable, limiting distribution of the probability distribution are well-known.

The set of options being advocated is a state, written as

$$S = (\theta_1, \theta_2, \dots, \theta_M) \tag{12}$$

Where θ denotes an option under consideration and M is the number of options being advocated in a particular case. M is often equal to N (the number of actors) but need not always be. Some actors may advocate the same position, setting $M < N$.

For the state S, the probability of option α in the limiting distribution is written as $P[\alpha \mid S]$. The entire process of determining the generalized votes, the coalitions, the transition probabilities and the limiting distribution is a PCE. We term the most likely option the 'probabilistic Condorcet winner' (PCW).

When the standard deterministic Condorcet winner (DCW; as described in *Unidimensional KTAB*, section 4.2, page 19) exists, it is almost always the PCW.

Under certain carefully designed sets of options and voting rules, it is possible to determine a closed form solution for the PCW. A closed form solution is one that can be written out algebraically through a series of equations; open form solutions require simulations to derive possible answers. One celebrated example is the Median Voter Theorem presented by Duncan Black (1948), using a deterministic special case of the general theory presented here.

2.2 Using Probabilistic Condorcet Elections to Model CDMPs

The PCE can be used to develop many models of bargaining and negotiation among agents. The general theory is as follows.

The expected utility of the state S to actor i is the following:

$$u_i(S) = \sum_{m=1}^{M} u_i(\theta_m) P[\theta_m | S] \tag{13}$$

The equation is really nothing more than a weighted average, calculated as the sum of the utility of each option ($u_i(\theta_m)$) multiplied by the probability of that option ($P[\theta_m \mid S]$).

This raises the first important point: the utility of any state for an actor is a trade-off between the utility of the option and its probability.

This enables us to model strategic voting, where actors look beyond the narrow utility of an individual option and instead vote for favorable outcomes. An actor might propose and vote for an option that has a lower utility to it than the preferred option but by virtue of its higher probability (an indication of support from other actors) could emerge as the winner and so block the emergence of an option with even lower utility.

How, then, does an actor 'decide' to change positon? Suppose that in state S, actor i changed its position from θ_i to α. By way of example, we can view this as actor 2 moving from

$$S = (\theta_1, \theta_2, \theta_3, \theta_4)$$

to

$$S' = (\theta_1, \alpha, \theta_3, \theta_4)$$

In shorthand, we can write this new state, S', as

$$S' = \sigma(i, \alpha \mid S) \tag{14}$$

where σ denotes a substitution.

The utility which actor i gains by changing its own position depends on not only that value to itself of that position, but also how likely that and the other positions are to occur. This is just the utility to i of the new state, S' which we can calculate in exactly the same way as for state S :

$$\begin{aligned} u_i(S') &= u_i\big(\sigma(i, \alpha \mid S)\big) \\ &= \sum_{m=1}^{M} u_i(\theta_m) P[\theta_m \end{aligned} \tag{15}$$

Thus, the gain in utility for actor i is the difference between the two utilities:

$$\begin{aligned} g_i(i, \alpha \mid S) &= u_i\big(\sigma(i, \alpha \mid S)\big) - u_i(S) \\ &= u_i(S') - u_i(S) \end{aligned} \tag{16}$$

A CDMP can be modeled as a series of steps. In each step, each actor independently selects its next position to maximize its gain g_i. When no actor changes its position, the CDMP stops: at this point the actors are in a Nash equilibrium where each has maximized its expected utility given the positions of the others.

3. The Fundamental Steps of Bilateral Negotiation

In *Unidimensional KTAB* we presented the outcomes of CDMPs as dependent upon simple voting rules that could be solved mathematically. In the section above we have again described how we can determine winning outcomes given the different coalitions of support.

In both cases, the mathematics summarized a complex story: that actors were putting forth proposals and counter-proposals and voting for and against the various options in accordance with their influence, the perceived stakes and the other variables we have introduced. We now explain what we really mean by this simplified explanation.

We begin any simulation with the data collected. As a snapshot of actor positions along with influence and salience scores, these represent a State. The initial state, determined by the collected data, is State 0. Subsequent turns of the simulation produce new states.

In the simple dynamic CDMP, the actors change their positions by generating a series of proposals and counter-proposals until a stable state is reached. Actors craft proposals that improve the likely outcome from their perspective, while winning enough support that the whole group prefers the new proposal to the old. They do not change their criteria of what is desirable; they change their positions to better achieve those criteria. This continues until no actor can improve their position, at which point a stable state has been reached. The most likely outcome of the simple dynamic CDMP is not necessarily one to which all actors agree: weak actors might be overruled by strong ones. The probable winner for each turn is shown through a probability curve.

In none of the models put forward in *Unidimensional KTAB* did actors change their positions or utility functions; actors merely had positions and voted accordingly. This paper describes how we add

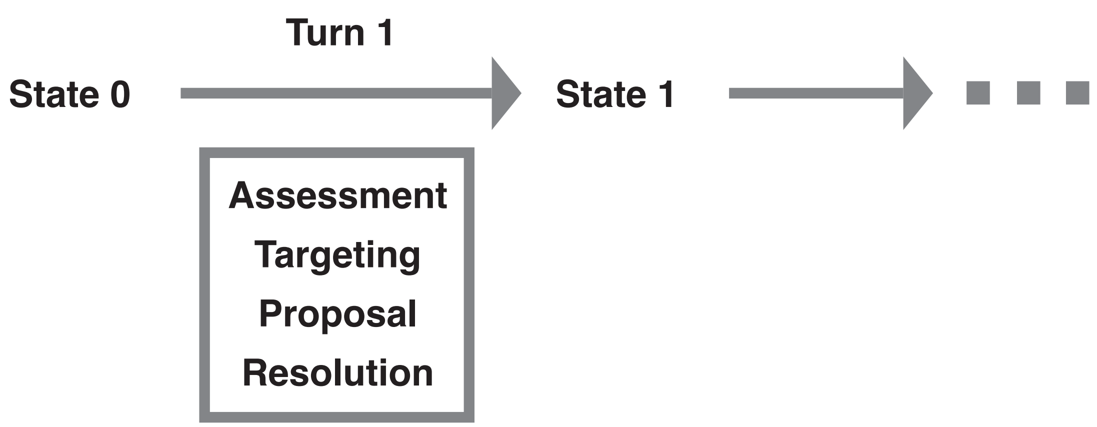

Figure 3.1 - States represent the set of positions being advocated, along with influence and other scores. The model is run during each turn to generate a new state. Source: KAPSARC

an extra level of sophistication in the modeling of how one actor tries to induce a change in position of another actor through a negotiation round. Inducement here simply refers to the ability of one actor to persuade another to change both its underlying position and view of utility. This is akin to changing someone's opinion to coincide with your view, rather than just getting them to support you without actually agreeing with you.

This is the complex dynamic CDMP presented in *Chinese SOE Reform*: a stylized negotiation simulation that includes not only voting according to some voting rules, but also negotiations in which one actor tries to induce another to change its position (by altering its view of its own utility curve).

Each negotiation turn can be understood through the formalized lens of a stylized model of four parts which, in each round of bilateral negotiations, must be performed in their set order. Note that while the four steps must be carried out in sequence, within each step the calculations for each actor can be made in parallel.

1. **Assessment.** Separately, each actor, i, develops its own perspective on the current situation. This includes estimating the 'revealed preferences' of other actors, based on their current positions.

2. **Targeting.** Separately, each actor, i, assesses all the other actors to identify another actor, j, whom it would be most beneficial to try to influence. This includes assessment of the stakes for each side, as well as the likely support or opposition of all third parties who might or might not become involved.

3. **Proposal.** If actor i can identify a worthwhile target j, then actor i develops a proposal of what new positions i and j should adopt. As j was selected to be weak compared with i, including in terms of support from third parties, this is likely to involve large concessions on j's part and only minor trade-offs on i's part. This bargain is developed as the Nash Bargaining Solution, taking into account the differing utilities and beliefs of the two actors.

4. **Resolution.** Generally, after the previous step, each actor will have several bargains which they either sent or received. Often, very weak actors will be targeted by multiple stronger actors. At most, one of the stronger actors can impose a new position; possibly none will succeed if the weak actor can rally enough allies to maintain its most preferred, status quo position. Mid strength actors may send bargains to those weaker than themselves, while simultaneously receiving bargains from those stronger than themselves, but at most only one of the bargains will actually occur because the actor can have only one new position. Separately for each actor, the competing bargains are resolved by an assessment of the balance of power favoring each alternative. The winner–possibly the status quo–is the actor's new position.

As a result of these many calculated interactions, there is a high degree of complication in the set of potential results that can emerge on a turn-by-turn basis. Consequently, a detailed narrative can emerge for each individual actor's behavior in each simulation.

With the underlying theory described and the basic concepts of negotiations established, we can now step through the logic of each stage.

4. The Mathematics of Bilateral Negotiation

4.1 Assessment

4.1.1 Assessing Distances

As mentioned earlier, a key feature of the SMP is that each actor regards its own current position as the most strategically desirable, with declining utility assigned to positions further away. Because actors attribute different levels of salience to different dimensions, each actor uses its own salience weighted Euclidean distance to assess the difference between another position and its own:

$$d_i^2(\theta) = \frac{\sum_k \left(s_{ik}\,(\theta_{ik} - \theta_k)\right)^2}{\sum_k (s_{ik})^2} \qquad (17)$$

where s_{ik} is the salience to the i-th actor of the k-th dimension, θ_{ik} is the coordinate of the presentation of the i-th actor over the k-th dimension (see equation 3).

In one dimension, this simply reduces to the distance between two points on a line:

$$d_i(\theta) = |\theta_i - \theta| \qquad (18)$$

where $d_i(\theta)$ is the distance of actor i (at position θ_i) from an arbitrary position θ.

The K-dimensional position is inside a K-dimensional box, $[0,1]^K$. Therefore the maximum difference along each axis is 1, and so the maximum possible salience weighted difference is also $d_i(\theta) = 1$. The minimum is zero difference, when $\theta_i = \theta$.

For convenience, we use the following abbreviation: $d_{ij} = d_i(\theta_j)$, where d_{ij} is the distance of actor i from the position adopted by actor j. Note that usually $d_{ij} \neq d_{ji}$ since because the different actors have different saliences, they measure distance differently.

In the SMP, all the actors must have the same understanding of the structure of the issue so all actors agree on the matrix d_{ij}. This is not a requirement in the CPT, as described in *Unidimensional KTAB*.

4.1.2 Assessing Risk Tolerance

An essential feature of actors is their attitude toward risk, which is determined by the shape of their utility curve. We will use the standard textbook definition of risk aversion, with three categories:

1. Risk neutral;
2. Risk averse; and
3. Risk seeking.

Figure 4.1 gives the utility curves for four different actors, Red, Green, Blue, and Brown broken line, against distance.

Risk neutral: Consider the green line. This indicates the utility to Green of other positions, which declines with increasing distance. The utility of a position at distance 0 (i.e. Green's own position) is exactly 1, the utility for a position at distance 0.2 is 0.8, and so on out to the utility of 0 for a position at distance 1. In general, Green's utility can be described as $u = 1 - d$. For clarity, we will color numbers denoting utility **blue** and those denoting distances **red**; those denoting probabilities will be shown in **grey**.

If Green were to be offered a risky situation where there would be 70% chance of achieving an outcome at distance 1 and 30% chance of achieving an outcome at distance 0, then the expected utility to Green of that risky situation is 0.3:

0.3 = 0.7 x (1 - 1) + 0.3 x (1 - 0)

Consider a riskless situation in which Green can achieve exactly the expected outcome, which has a distance of:

0.7 = 0.7 x 1 + 0.3 x 0

The utility to Green of this riskless situation is therefore:

0.3 = 1 - 0.7

Thus, because the green utility curve is precisely a straight line, the average utility of the outcome is the utility of the average outcome. Green has no preference between a risky and riskless situation, so long as the expected outcome is the same. Further reasoning along these lines shows that Green would make choices between situations based solely on the expected outcome, regardless of the risk inherent in each; such behavior is called 'risk neutral'.

Risk averse: Consider the red and blue utility curves, which are not straight lines. For both Red and Blue, the utilities of the two possible outcomes are the same as they were for Green: utility 0 for the outcome at distance 1 and utility 1 for the outcome at distance 0.

Looking at the risk case outlined above, the average utility of the outcome is still 0.3 for Red. However,

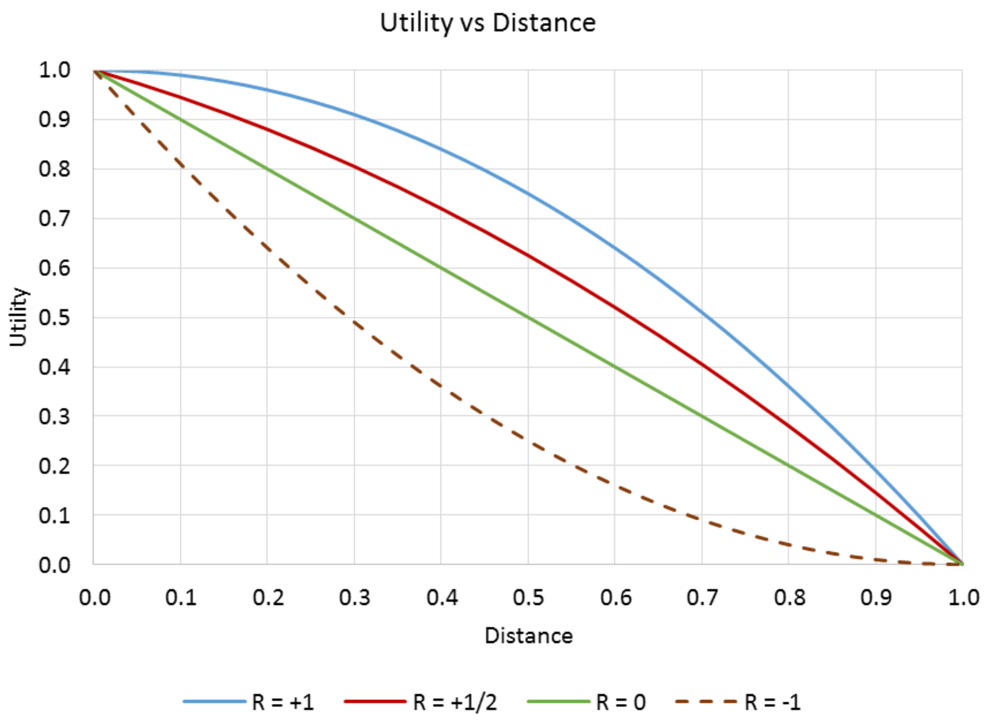

Figure 4.1 - Curved Utility Functions – Four different risk tolerances. Green (risk neutral): utility falls linearly with distance such that the utility of the average is the same as the average utility making Green neutral between them. Blue and Red (risk averse): utility curve is convex, with utility falling only a little over short distances but more over long distances. The average utility (risky) is lower than the utility of the average (riskless), and therefore a risk premium is demanded. Broken Brown (risk seeking): utility curve is concave.

Source: KAPSARC

utility of the average outcome (0.7) is approximately 0.4. For Red, the utility of the average outcome is greater than the average utility of the outcomes. Red would have to be paid a risk premium to take the risky situation over the riskless situation, even though both have the same expected outcome (a distance of 0.7). This is the textbook definition of 'risk averse': the actor has a convex utility curve, and hence demands a positive risk premium.

Risk averse behavior is taken for granted throughout the world of economics and finance, in which the risk free assets offer the lowest interest rates, and progressively more risky assets have to offer progressively higher risk premiums. Indeed, increasing levels of interest on loans-compared with the risk free rate-are often taken as a surrogate for market perceptions of risk of default, precisely because all the financial actors are risk averse, to one degree or another.

Because Blue's utility curve has greater curvature, Blue would demand an even higher risk premium than would Red. In other words, greater curvature corresponds to greater aversion to risk, so higher risk premiums are demanded to make up for it. With this parameterization, Blue's risk premium is exactly twice Red's risk premium.

Risk seeking: in theory, actors could be risk seeking, which would appear as a concave utility curve (similar to the Brown broken line). This is sometimes observed for trivial risks, such as gambling small amounts of money simply for the excitement, while knowing that 'the house always wins'. Such behavior rarely persists when the stakes are significant. For example, literally risk seeking investors would seek out risk as being a good thing in itself, so they would voluntarily choose a financial instrument which was guaranteed to lose money, as long as it made up for that by running a high enough risk of losing even more money. Needless to say, markets would be swift to take advantage of such behavior, with serious adverse consequences for the risk seeking investors.

The demonstration version of **smpc** currently in KTAB includes the risk tolerance parameter R to describe the curvature of actors' utility curves. These can range from very risk tolerant, to risk neutral, to highly risk averse, so that $0 \leq R \leq 1$, as per Figure 4.1. Of course, some model builders might choose to use negative R values for risk seeking if they decided that was appropriate.

Compromise: in a bargaining context, curvature also manifests itself as willingness to compromise. Consider the changes in utility for Red, Green and Blue from a very small compromise, say to a new position at a distance of 0.1. Green would perceive this as involving a 0.10 decline in utility, Red would perceive this as costing about 0.05, while Blue would perceive it as costing about 0.01. Thus, Blue would be the most willing to bargain away 0.1 of its position, while Green would be the most resistant to doing so. This behavior has nothing to do with uncertainty or risk, as the terms of the bargain are precise and deterministic: a shift by 0.1.

This compromise behavior occurs whether the bargains are voluntary or coerced. It would take a small promise of reward to make Blue shift its position (voluntary) or only a small threat of harm (coerced). It would take the largest reward to make Green shift (voluntary) or the largest threat (coerced).

Given the salience-weighted distance function and the values, the utility curve for each actor can take a particularly simple form:

$$U = (1 - d)(1 + Rd) \qquad (19)$$

If we include all subscripts and parameters, the same formula appears as follows.

$$U_i(\theta) = (1 - d_i(\theta))\,(1 + R_i\, d_i(\theta)) \qquad (20)$$

This is the equation that has been used to draw Figure 4.1 and gives the following:

- In the risk neutral case $R = 0$, this reduces to the straight green line: $U = 1 - d$;
- In the risk averse case $R = +1$, this reduces to the convex blue parabola: $U = 1 - d^2$; and
- In the risk seeking case $R = -1$, this reduces to the concave brown parabola: $U = (1 - d)^2$.

This utility function is particularly simple, but it is confined to the differences and utilities in the range [0,1].

In order to bargain (or even to determine which other actors are good targets for bargaining), it is necessary to have an idea of their relative willingness to bargain, i.e. to estimate their parameters. In the current demonstration **smpc** model, this is done by looking at R in terms of revealed risk aversion.

As mentioned earlier, the PCE produces an estimate of the probability that each actor's position will actually be adopted. Analogous to 'revealed preference' in economics, the willingness of actors to advocate positions which have little chance of success can be seen as revealed risk tolerance. Conversely, those actors which have adopted the positions most likely to succeed have revealed a high aversion to risk. This reasoning is reflected in the following algorithm.

To estimate the R_i parameters for each actor, we first perform a PCE, which determines the probability for each actor's position. Probabilities are assigned to the position, not to the actor, because several actors might adopt the same position. If that position is adopted, then each one of the actors has its position adopted, because they are all the same position. Thus, the sum of probabilities over positions is exactly 1, while the sum of probabilities over actors is often greater than 1 because of double counting (or triple, or quadruple, or more).

The highest and lowest probabilities are denoted as P_H and P_L respectively. We rescale the probabilities so that the most likely actors are assigned a risk aversion of 1 and the least likely actors are assigned a risk aversion of 0:

$$R_i = \frac{P_i - P_L}{P_H - P_L} \qquad (21)$$

If the modeler wishes to use in the ranges [-1, +1] or [-1/2,+1], similar formulae for R can easily be derived.

The alert reader will note that there is a circular dependency in this reasoning. The R_i values are determined from P_i values, which are determined by the PCE, which depends on utility values, which depend in turn on the R_i values. This circular dependency is broken by employing a two-step method.

1. R_i values are set to zero and a PCE conducted over the positions; then
2. The calculated P_i values are used to determine the second set R_i values.

Obviously, the two sets of R_i values differ, so this is in some sense not a self consistent estimate. This process could be iterated so as to arrive at a stable and self consistent set of R_i values, but this has been numerically observed to have little effect on the final R_i values. Further, the difference between the two-step estimate and the self consistent estimate are small compared with the effects of anchoring and adjustment (see next section). The two-step method is therefore a useful approximation.

Alternative utility functions in KTAB

As KTAB is a toolkit for building models, it is worth a short diversion to consider when alternative utility functions might be more useful. In economics and finance, a common utility function which is not limited to [0,1] is the negative exponential, where a is the risk aversion parameter:

$$U(x) = \frac{(1 - e^{-ax})}{a}$$

Note that when $|ax|$ is small compared with 1, the utility curve is approximately linear and risk neutral: $U(x) \approx x$. Risk aversion exerts more of an effect as the potential gains or losses grow larger.

If currency is measured in USD, then the units of a are 1/USD, so the right hand side of the equation has units USD, as required. For example, if the actor saw losing amount D as being twice as harmful as gaining D is helpful, then its risk aversion parameter would be $(ln2)/D$. The investor's behavior would be approximately risk neutral as long as the potential gains or losses x were small compared with $D/(ln2)$ so that $|ax|$ would be small compared with 1.

For any risky asset, the 'certainty equivalent', C_x, is that dollar amount for which a risk averse investor would assign equal utility to the certain dollar amount and to the risky asset. Consider an asset whose return x is normally distributed with mean μx and variance σ_x^2. With the negative exponential utility function, the certainty equivalent of a normally distributed return takes the familiar mean variance form:

$$C_x = \mu_x - \left(\frac{a}{4}\right)\sigma_x^2$$

Note that for a risk seeking investor, $a < 0$, so no matter how negative is the expected return, a large enough risk will render the investment attractive.

For models of economic and financial bargaining, a KTAB modeler could choose to rescale outcomes into a [0,1] utility range, or might choose to work directly with the more familiar C_x values. A procedure for estimating a would have to be developed, such as estimating an actor's revealed risk attitude by finding that a which maximizes the certainty equivalent of the actor's initial position. In either case, the basic bilateral negotiation process could still be employed, as it is not dependent on the particular shape of any actor's utility functions.

4.1.3 Adjusting Risk Tolerances

The phenomenon of ‘anchoring and adjustment’ is widely observed in psychological studies and is routinely exploited in efforts to manipulate opinion. The general statement is that people estimate quantities by extrapolating from what is familiar, but they usually do not extrapolate enough.

This is reflected in KTAB by differentiating between an actor’s actual risk attitude and how other actors estimate their own risk attitude.

We denote actor h’s estimate of actor i’s risk attitude by R_i^h If the actors remain fully anchored and do not adjust, then their estimate of another’s risk attitude is exactly their own:

$$R_i^h = R_h^h$$

While this might seem like an implausibly bad estimator, the phenomenon is common enough that strategic planners are routinely cautioned against ‘mirror-imaging’, i.e. assuming that their opponents have the same utility values, and hence risk tolerance, as themselves.

Another estimator is simply that actors fully and completely adjust their estimates, and always perceive the other actors perfectly:

$$R_i^h = R_i^i$$

Given the ubiquity of anchoring and adjustment, a more plausible estimator might be one where the actors take their own risk attitude as the anchor, and adjust somewhat toward the other actor’s true attitude, such as the following, where the estimator is weighted two-thirds anchored, one-third adjusted:

$$R_i^h = \frac{2R_h^h + R_i^i}{3} \qquad (22)$$

Whichever version of adjustment is selected, this leads immediately to actor specific estimates of all the relevant quantities, such as utilities:

$$U_i^h(\theta) = (1 - d_i(\theta))(1 + R_i^h d_i(\theta)) \qquad (23)$$

These utilities can be used in an actor specific PCE to also determine P_i^h for that actor’s subjective probabilities. The actor specific numbers are used extensively in the second step of our simulation of bilateral negotiations, targeting.

While the details of bargaining are discussed later, we note now that the variation among actor specific estimates can play a very significant role in bargaining. For example, differing estimates might lead two actors each to think it is likely to win if it were to initiate a conflict with the other, so each one goes into the bargaining thinking it has the upper hand. Conversely, actors A and B might both believe that A has the upper hand over B, but they disagree as to how much. If A thinks it has only a small advantage, while B thinks that A has a large advantage, then A might seek much smaller concessions than could have been obtained.

While the specific amount of adjustment (e.g. $\frac{1}{3}$, $\frac{1}{2}$ or $\frac{2}{3}$ of the difference) does change the precise magnitude of such differences, it does not change the signs or the qualitative structure

Notice that in the basic SMP, KTAB supports nine kinds of assessment models, depending on the following parameterization. The choices in the demonstration **smpc** can be changed simply by switching an enumerated option (default settings are in italics):

- R-range: [*0,1*], or [-1/2, 1], or [-1,+1]
- R-adjustment: full, *partial*, or none

4.2 Targeting

The essence of the targeting step is that each actor, i, tries to identify the other actor, j, over which i has the largest advantage and thus can most likely shift to a position more to i's liking.

This analysis is done by analyzing not only a game tree between the two principal actors i and j, but also some subsidiary game trees for third parties. To aid in formalizing these concepts, we introduce three new notations:

> The symbol $i \rightarrow j$ denotes the event in which i tries to change the position of actor j. Actor might simply acquiesce if the issue has low salience, or might resist.

> The symbol $i : j$ denotes a contest between actors i and j as they try to change each other's position, each backed by a coalition of third-party supporters.

> The symbol $i \succ j$ denotes the event that i successfully imposes its position on j, and vice versa for $j \succ i$.

- This is generally a hypothetical event, where each actor analyzes the likely consequences of an outright struggle, so as to determine what concessions can be plausibly obtained in bargaining.
- The curved $\succ$ symbol is distinct from the straight numerical inequality $>$ symbol .

The set of positions currently advocated by all the actors at the start of the turn is termed the 'state' at that point in time. Each actor assigns a value to the entire state, which is simply the average value of all the actor's positions.

Unfortunately, at this point we must introduce a plethora of subscripts and superscripts to express the many possible combinations that will arise.

4.2.1 Utilities

Remember that the utility of a state – $U(S)$ – is equal to the expected utility of all the individual positions – θ_k. In this instance we are assuming equal probability of each position, hence dividing by the number of positions, $\frac{1}{M}$, in the status quo ante.

Actor h's estimate of the utility to actor n of the state is simply the average over the initial positions of all N actors:

$$U_n^h(S) = \frac{1}{M}\sum_{k=1}^{M} U_n^h(\theta_k) \tag{24}$$

The utility of the hypothetical state in which i successfully gets j to adopt i's position is therefore the following:

$$\begin{aligned} U_n^h(\, i \succ j|S) &= U_n^h[\sigma(\, j, \theta_i|S)] \\ &= \frac{1}{M}\sum_{k=1}^{M} U_n^h(\theta'_k) \end{aligned} \tag{25}$$

Where all the positions have a uniform probability and are unchanged except that j's position is replaced by i's position.

When the meaning is clear from context, we will sometimes use the following compact notation:

$$U_{ij}^h = U_i^h(\theta_j) \tag{26}$$

and

$$U_{n,ij}^h = U_n^h(i \succ j|S) \tag{27}$$

Thus, h's estimate of the gain to n from i defeating j would be

$$U_{n,ij}^h - U_n^h(S).$$

We could expand all the summations and show that most of the terms cancel in this subtraction; later on it will become clear that the $\frac{1}{M}$ terms cancel when computing probabilities. Nevertheless, we present the formulae in this way because it makes clear that evaluation over whole states is involved, and in the Resolution step, the cancellations are not so neat.

The evaluation of an entire state is crucial to the representation of naïve versus strategic voting.

> Naïve voters simply vote for whatever policy which they expect would have the best outcome if implemented. As mentioned earlier, evenly weighting the positions in a state is essentially the same as voting according to the value of positions: it models naïve voting.

> Strategic voters allocate their vote so as to get achieve the best outcome for themselves, given how they expect other voters to behave. This can lead to much more sophisticated behavior, such as supporting proposals not because they are desirable in themselves but because they help split the support for a proposal seen which was even less desirable, thus allowing a more desirable alternative to win.

Using an embedded PCE to estimate the probability of each outcome, rather than just assuming equal probabilities, has the effect of modeling strategic voting, because it allows the actors to estimate how their actions affect the probability that second or third parties will succeed or fail.

Although the differences between naïve and strategic voting are quite profound, all the following formulae should be taken as equally applicable to both cases, because they differ only in how the utility of a whole state is estimated. This is just one example of how KTAB provides a unified framework for modeling quite diverse behaviors.

As mentioned earlier, the total salience of an actor is essentially the probability that it will resist changes to its position, rather than acquiesce because it is focused on other (un-modeled) issues. Thus, the expected value of a challenge is the probability-weighted average of the value of acquiescence and a contest:

$$U_n^h(i \Rightarrow j \mid \mathrm{S}) = (1 - s_j) U_n^h(i \succ j|S) + s_j U_n^h(i{:}\, j \mid \mathrm{S}) \tag{28}$$

In turn, the expected utility of a contest is the probability-weighted average of the value of success or failure:

$$U_n^h(i{:}\, j \mid \mathrm{S}) = P^h(i \succ j|S) U_n^h(i \succ j|S) + P^h(j \succ i|S) U_n^h(j \succ i|S) \tag{29}$$

The criterion for i's selecting the best target j is to maximize i's expected gain to itself of challenging j, compared with the status quo:

$$U_i^i(i \Rightarrow j \mid \mathrm{S}) - U_i^i(S) \tag{30}$$

If that maximum is zero or negative, then i is too weak to try to change any other actor's position.

4.2.2 Generalized Voting

While generalized voting has been mentioned several times earlier, at this point it is necessary to formalize the concept.

Suppose actor i is faced with a choice between two alternatives, x and y. i's generalized vote, $v_i(x{:}y)$ is the amount of influence exerted to promote x over y. It is positive if i favors x; it is negative if i favors y. These alternatives might be the positions of individual actors, but they might be more complex alternatives (such as a set of positions, or state), so we do not denote the alternatives simply by θ.

Note that it is quite likely that $v_i^h(x{:}\, y) \neq v_i^i(x{:}\, y)$ because the actors' estimated utilities usually differ.

The KTAB toolkit offers five different rules for generalized voting. The main **smpc** model assumes that each actor i has a single capability measure for maximum overall influence, c_i. Because actors have different overall saliences, they will exert that capability to differing degrees; for brevity we employ the notation $w_i = s_i c_i$. Again, we require $s_i < 1$ so that the influence actually exerted, w_i , does not exceed the maximum capability to exert w_i, does not exceed the maximum capability to exert influence, c_i. Note that the different saliences of different dimensions are used in the calculation of the utility by way of the distance.

1. **Binary voting** is the exertion of all the actor's influence to promote whichever outcome it prefers:

$$vb_i^h(x{:}\,y) = \begin{matrix} +w_i, & U_i^h(x) > U_i^h(y) \\ -w_i, & U_i^h(x) < U_i^h(y) \end{matrix} \tag{31}$$

2. **Proportional voting** is the exertion of the actor's influence in proportion to its perceived stakes, which is how much it stands to gain or lose from the two alternatives:

$$vp_i^h(x{:}\,y) = w_i\left[U_i^h(x) - U_i^h(y)\right] \tag{32}$$

Note that in the SMP, utilities are always in the [0,1] range so that the difference in utility is always in the [-1,+1] range, which avoids artificially extending an actor's influence beyond the [$-w_i$,+ w_i] range.

3. **Cubic voting** is designed to model the situation where actors respond little to small changes, but respond energetically to large losses and to large gains. The motivation for including this voting rule comes from the political advice "focus benefits and diffuse costs". The point of the advice is to exploit situations where a few highly motivated beneficiaries will exert more total influence than the many, but barely motivated, cost-bearers.

$$vc_i^h(x{:}\,y) = w_i\left[U_i^h(x) - U_i^h(y)\right]^3 \tag{33}$$

Again, note that with utilities in the [0,1] range, this avoids artificially extending an actor's influence beyond the [$-w_i$,+ w_i] range.

For w_i = 1, the three voting rules are contrasted in Figure 4.2.

Currently, KTAB also offers two more voting rules which introduce moderate nonlinearity into the proportional voting rule:

4. **Binary-Proportional voting** mixes proportional and binary voting. This results in a linearly sloped response curve, but with a (nearly) vertical jump around 0:

$$vbp_i^h(x{:}\,y) = (1-\lambda_b)vp_i^h(x{:}\,y) + \lambda_b vb_i^h(x{:}\,y) \tag{34}$$

The slope factor λ_b=0.2 was selected so that the vertical jump in the middle was half the height of either linear slope. For w_i=1, the maximum vertical deviation from the Proportional voting rule is $\pm\lambda_b$ right at zero utility difference.

5. **Cubic-Proportional voting** mixes proportional and cubic voting. It is approximately linear in the mid -range, but with curves up and down at the extremes:

$$vpc_i^h(x{:}\,y) = (1-\lambda_c)vp_i^h(x{:}\,y) + \lambda_c vc_i^h(x{:}\,y) \tag{35}$$

Again, for w_i = 1, the VPC rule is approximately linear around zero and passes through the (0,0) origin. Although the rules coincide at zero, the Proportional-

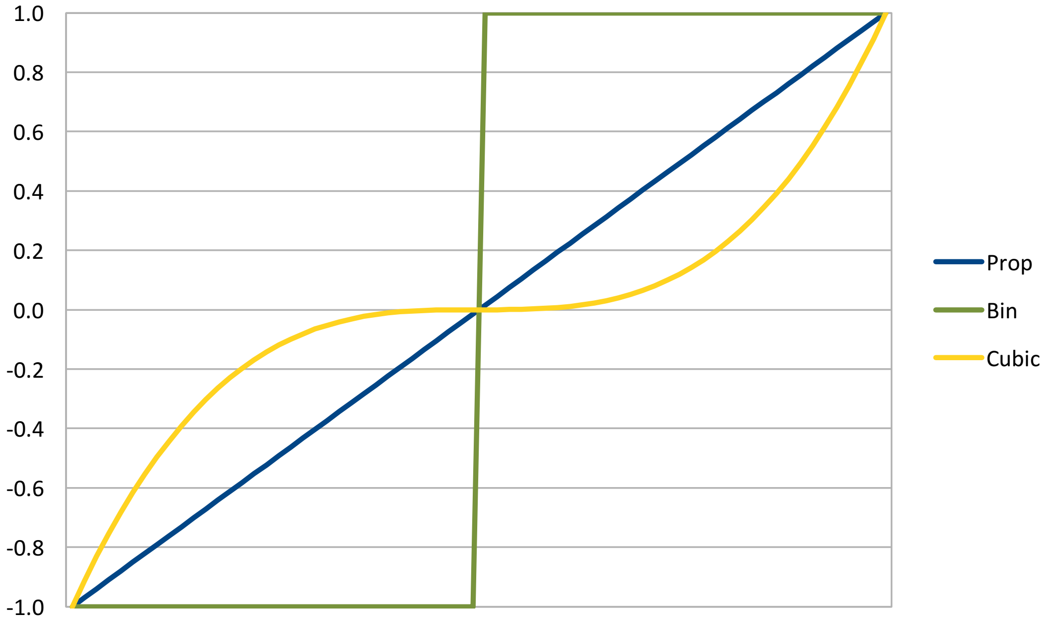

Figure 4.2 - Proportional, Binary, and Cubic Voting Rules. The amount of effort an actor exerts, compared to the stakes, is described by a voting rule. The simplest is the proportional voting rule, as depicted here in blue: the amount of effort depends directly on the stakes. Another rule is the binary (either/or) response, depicted here in green: the actor always exerts the maximum effort they can for their favored option, no matter how low the stakes. Another rule is the cubic response, depicted here in yellow: actors hardly react at all to small stakes, but their response then grows dramatically as the stakes increase.

Source: KAPSARC

Cubic rule has a slope of only 1- λ_c around the origin, compared with the Proportional rule's slope of exactly 1. The maximum vertical deviation occurs at two distinct points on either side of zero. If we set λ_c as follows, the maximum vertical deviation of the Cubic-Proportional rule (on either side of zero) will be the same as the maximum vertical deviation of the Binary-Proportional rule (exactly at zero):

$$\lambda_c = \frac{3\sqrt{3}}{2} \lambda_b \qquad (36)$$

With λ_b = 0.2, we get λ_c = 0.5196

For this parameterization, with w_i = 1, the Proportional, Binary-Proportional, and Cubic-Proportional voting rules are contrasted in Figure 4.3.

When we wish to indicate that any of the five voting rules could be used in voting between alternatives x and y, we use the basic symbol $v(w,U_x,U_y)$ or $v(x{:}y)$ for short. This is the usual case: wherever a voting rule is used, any of them could be used.

4.2.3 Coalitions and Probabilities

While the utility and salience terms $U_i^i(i \rightarrow j \mid S)$ have been described, the probability estimates have not. The probability that the the i-th actor will succeed in changing j's position depends critically on the strength of each coalition backing i or j. Thus, a comparatively weak actor with many strong allies could easily pressure a comparatively strong but isolated actor into making concessions. Note that 'strength' is here a mixture of both intrinsic capability and the perceived stakes, which in turn depend on both policy distance and salience.

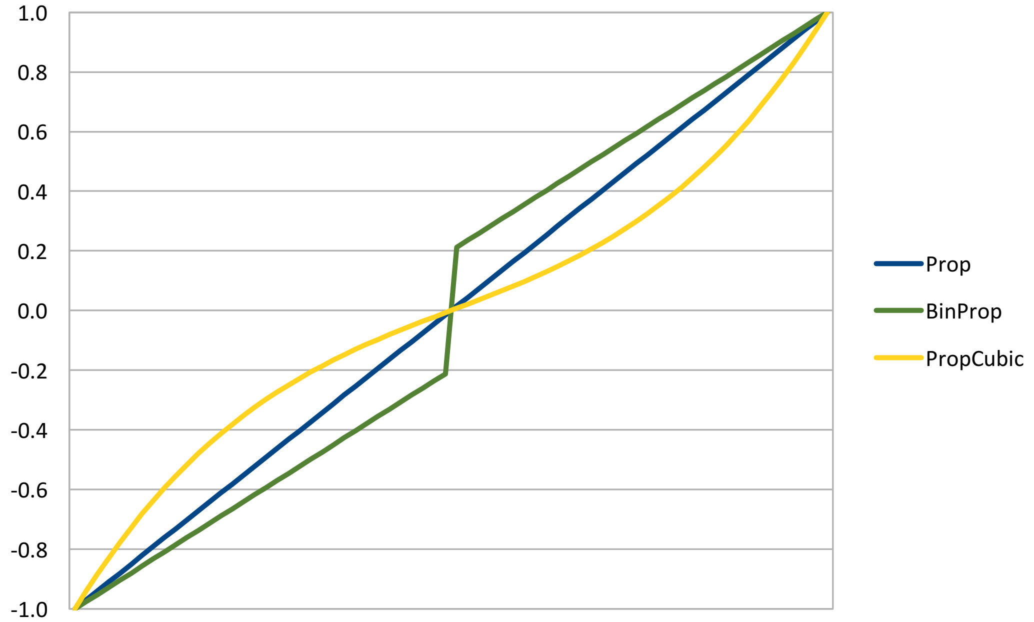

Figure 4.3 - Proportional, Binary-Proportional, Cubic-Proportional Voting Rules. The amount of effort an actor exerts, compared to the stakes, is described by a voting rule. The simplest is the proportional voting rule, depicted here by the blue diagonal: the amount of effort depends directly on the stakes. The binary and cubic rules from the previous figure have been contracted toward the diagonal so that their deviations from proportionality are less extreme. The green binary-proportional rule keeps the discontinuous jump from favoring to opposing, but it now has some gradation in the response after that. The yellow proportional-cubic rule keeps the reduced response for small stakes, but it now has some slope even near zero stakes.

Source: KAPSARC

Please refer back to section 2.1. Remember that the strength of the coalition backing alternative x over alternative y is just the sum of the influence which various actors exert in favor of x:

$$c^h(x:y) = \{n | v_n^h(x:y) > 0\} \tag{37}$$

$$s^h(x:y) = \sum_{n \in c^h(x:y)} v_n^h(x:y) \tag{38}$$

In section 2.1 we also stated that the probability that option x is selected over option y is simply:

$$P^h(x \succ y) = \frac{s^h(x:y)}{s^h(x:y) + s^h(y:x)} \tag{39}$$

We can now be more precise in our discussion of how the probabilities are calculated. Currently, KTAB provides two formulae for estimating the probability that one coalition will prevail over another.

1. The first assumes that probabilities vary linearly with coalition strength:

$$PS^h(x \succ y) = \frac{s^h(x:y)^2}{s^h(x:y)^2 + s^h(y:x)^2} \tag{40}$$

And similarly for $PL^h\ (y \succ x)$.

Under the linear law, a 2:1 advantage in strength gives a 67% chance of success. Many empirical studies of strength (strength of sports teams, strength of political coalitions, etc.) assume the linear law and then use the observed probabilities to estimate the implied strengths.

2. The second rule treats probabilities as varying with the square of coalition strength:

$$PS^h(x \succ y) = \frac{s^h(x:y)^2}{s^h(x:y)^2 + s^h(y:x)^2} \quad (41)$$

And similarly for $PS^h\ (y \succ x)$.

Under the square law, a 2:1 advantage in strength gives a 80% chance of success. Some empirical studies of the strength of military units take the observed probabilities and then use either PL or PS to estimate the implied strengths, depending on whether the analyst chooses the Lanchester Linear Law (more commonly used with ancient warfare) or the Lanchester Square Law model of combat (more commonly used with modern battles).

When we wish to indicate that either rule may be used, we use the basic symbol P.

The alert reader will notice that, in terms of statistical estimation from known probabilities, there is no essential difference between these two models. Suppose one uses PS to estimate strengths $s_{i,S}$ and uses PL to estimate strengths $s_{i,L}$. Then $(s_{i,S})^2 = s_{i,L}$. However, in terms of estimating probabilities from strengths, there is a difference.

Remember that the strength of the coalition favoring x over y is the sum of the individual actors' support for x over y. The strength of the coalition favoring option x over x itself is trivially zero, and so are all the diagonal entries of the $C_{i:j}$ matrix used to tabulate the individual actors' support for x over y in the KTAB output. When probabilities are then calculated from these numbers, the diagonal entries of the probability matrix found in KTAB's output are set to 0.5. This is done for convenience to simplify the mathematics, rather like the decision that 0! is 1, rather than 0 or undefined.

When there are only two options, the limiting probabilities of the Markov Process, i.e. the stable PCE distribution, are exactly the probabilities originally determined by the strength ratios.

A simple example may help explain. Imagine seven actors choosing between two options using the proportional voting rule as per Table 4.1. We can then calculate the (linear) probabilities for each option straight from the initial coalition strengths:

$$0.42 = \frac{6.466}{6.466 + 8.856}$$

$$0.58 = \frac{8.856}{6.466 + 8.856}$$

The final probabilities are the simple ratio. However, with a PCE over three options, there is no simple formula for the stable distribution.

4.2.4 Principal Actor Voting

When actor i analyzes a possible effort to change the position of actor j, these are termed the principal actors and all others are termed third parties. While we have discussed the voting rules each might use, there are several different ways to assess the utilities they might use in voting.

While it might seem odd to analyze how actor i might vote in its own contest with j, some non-trivial interactions must be considered. For example, with the proportional voting rule, even when both actors see a large difference in position, the salience of the actors can have a large effect. There have been several instances in international relations where a powerful actor initiated conflict with a weaker actor, only to end up acquiescing because it (despite being more powerful) placed less salience on foreign adventures than the weaker actor placed upon national survival.

The simplest model of the two principal actors is that they use naïve voting with proportional voting. Because they could use naïve or strategic voting for each of the five voting rules, this range of parameterizations allows ten submodels of how the two principal actors behave. For all ten models of i's exertion of influence, the relevant difference is between i's estimate of the utility to i of i imposing its position on j, and i's estimate of the utility to i of j imposing its position on i:

$$v_i^i(i:j) = v(s_i c_i,\ U_{i,ij}^i,\ U_{i,ji}^i) \quad (42)$$

This will be positive, so it will contribute to i's estimate of the strength of the $c^i(i:j)$ coalition.

Similarly for i's estimate of j's vote:

$$v_j^i(i:j) = v(s_j c_j,\ U_{j,ij}^i,\ U_{j,ji}^i) \quad (43)$$

This will be negative, so it will contribute to i's estimate of the strength of the $c^i(j:i)$ coalition.

Actor	**Capability *w***	**Utility of option 1**	**Utility of option 2**	***vp*(1:2)**	***vp*(2:1)**
1	101.6	0.761	0.755	0.610	
2	155.6	0.739	0.792		8.246
3	104.7	0.947	0.946	0.105	
4	67.6	0.662	0.647	1.014	
5	41.6	0.776	0.737	1.621	
6	194.8	0.787	0.771	3.117	
7	50.8	0.626	0.638		0.609
Coalition strength				**6.466**	**8.856**

Table 4.1 - Calculating coalition strength. Seven actors choose between one of two options. The actors capability is shown, along with the utility each actor derives from each of the two options. The exertion of influence is then calculated (in this example) according to the proportional voting rule (see section 4.2.2). The strength of each coalition (with one coalition favoring option 1 over option 2, and the other coalition favoring option 2 over option 1) is simply the sum of the individual votes of each actor. Source: KAPSARC

4.2.5 Third Party Voting

As mentioned earlier, weaker actors can overcome stronger actors because of a difference in salience. Another way this can happen is through the influence of third parties, particularly strong third parties. Indeed, many of the conflicts during the Cold War were between comparatively weak actors, where each was backed by a coalition that included one super power on each side.

The simplest version of third party voting is naïve voting, where the only difference between $i \succ j$ and $j \succ i$ are the positions of the principal actors. Historically, many weak actors have backed a principal actor that they thought would win, expecting to gain wealth and power as part of a victorious coalition, only to find themselves suffering adverse consequences when their coalition was defeated. Similarly, strong third-party actors have intervened in what they thought were minor disputes, only to find themselves becoming committed to win or lose as part of an entire coalition.

These considerations lead to (at least) three submodels (see Tables 4.2-4.4) of how third parties, k, assess their stakes and hence how much influence to exert in the $i : j$ contest. We introduce the notation that $ik \succ j$ denotes the event where k supports i, then i and k combined defeat j; similarly for the other three possibilities. Regardless of k's behavior, the losing principal actor must (by definition of losing) adopt the position of the winning principal actor.

For any of these third party commitment rules, the value of the four resulting states can be computed as either the evenly weighted average (for naïve voting) or from a PCE (strategic voting). Each of the six possible submodels yields the four utilities $U_k^h(ik \succ j), U_k^h(j \succ ik)$, $U_k^h(i \succ jk)$ and $U_k^h(jk \succ i)$

Thus, the value to k of allying with i is the expected value of doing so:

$$U_k^h(ik:j) = P^h(ik \succ j)U_k^h(ik \succ j) + P^h(j \succ ik)U_k^h(j \succ ik) \qquad (44)$$

And similarly for $U_k^h(i:jk)$.

With those expected utilities in hand, the contribution of k follows directly from the voting rule:

$$v_k^h(i:j) = v(s_k c_k,\ U_k^h(ik:j), U_k^h(i:jk)) \qquad (45)$$

To analyze the value to k of the $(ik:j)$ alliance versus the $(i:kj)$ alliance in this restricted subgame, we must analyze not only the utility to k of the potential outcomes but also the probability of each potential outcome. This requires another assessment of mini coalitions supporting each side. To simplify the formulae, we define the following notation, which simply picks out positive values:

$$[x]^+ = \begin{cases} x, \text{if } x > 0 \\ 0, \text{otherwise} \end{cases} \qquad (46)$$

We consider only the restricted coalitions of i, j, and k. Thus, the strength of each coalition consists only of the vote of the principal actors, and a vote by k depending on which principal's position is preferred, and by how much:

$$s^h(ik:j) = v_i^h(i:j) + \left[v_k^h(ik:j)\right]^+ \qquad (47)$$

Similarly for the opposing restricted coalition:

$$s^h(j:ik) = v_j^h(j:i) + \left[v_k^h(j:ik)\right]^+ \qquad (48)$$

Uncommitted	θ'_i	θ'_j	θ'_k
$ik \succ j$	θ_i	θ_i	θ_k
$j \succ ik$	θ_j	θ_j	θ_k
$i \succ jk$	θ_i	θ_i	θ_k
$jk \succ i$	θ_j	θ_j	θ_k

Table 4.2 - Uncommitted. Third parties do not change their position, no matter who they support or who wins. This can be represented as a table of outcomes against the new positions of actors.

Source: KAPSARC

Semi-Committed	θ'_i	θ'_j	θ'_k
$ik \succ j$	θ_i	θ_i	θ_k
$j \succ ik$	θ_j	θ_j	θ_j
$i \succ jk$	θ_i	θ_i	θ_i
$jk \succ i$	θ_j	θ_j	θ_k

Table 4.3 - Semi-Committed. If its alliance wins, the third party can keep its initial, preferred position. If its ally loses, then it is forced to adopt the position of the winner, as follows.

Source: KAPSARC

4.2.5 Third Party Voting

Fully Committed	θ'_i	θ'_j	θ'_k
$ik \succ j$	θ_i	θ_i	θ_i
$j \succ ik$	θ_j	θ_j	θ_j
$i \succ jk$	θ_i	θ_i	θ_i
$jk \succ i$	θ_j	θ_j	θ_j

Table 4.4 - Fully Committed. Whichever side it allies with, it ends up having to adopt the position of the winning principal actor. Source: KAPSARC

With these restricted coalitions, the success probability is calculated as before:

$$P^h(ik \succ j) = \frac{s^h(ik{:}j)}{s^h(ik{:}j) + s^h(j{:}ik)} \qquad (49)$$

During the search for the best j to target, all calculations are done from i's perspective, so the estimator $h = i$. In the Targeting and Proposal phase, the perspectives of both i and j will be used, while in the Resolution phase, every actor's actions are determined by its own individual perspective.

Notice that even in the basic SMP, KTAB supports six hundred submodels of targeting, depending on how utilities, coalitions and voting are parameterized. The choices made in the current **smpc** demonstration can be easily changed by flipping a parameter value; the values currently used are italicized.

- Principal actor voting assessment: *naïve* or strategic;
- Principal actor voting rule: binary, *proportional*, cubic, proportional-binary, proportional-cubic;
- Third Party Commitment: full, *semi*, or none;
- Third party voting assessment: *naïve* or strategic;
- Third party voting rule: binary, *proportional*, cubic, proportional-binary, proportional-cubic; and
- Probability: *linear* or square.

4.3 Proposal

Aside from the distance matrix d_{ij}, the machinery of bilateral negotiations described so far is applicable to both SMP and non-SMP models. However, the process of developing proposed bargains in the SMP is heavily reliant on the fact that the SMP uses vector positions in K-dimensional space, for which continuous interpolation between positions is possible.

The fundamental approach to developing proposals is to identify advantageous targets from i's perspective, then estimate the result of bargaining between the two actors, given their differing perspectives and estimates. The fundamental tool for this is Nash Bargaining.

4.3.1 Nash Bargaining

John Nash (1953) presented a solution to the problem of bargaining between two agents. The essential result is that the expected bargain (B) maximizes the product of each actor's improvement over a reference outcome (R), where both terms of the 'Nash product' (η) are non-negative:

$$\begin{aligned} \eta_{ij}^{h}(B|R) &= \left(U_i^h(B) - U_i^h(R)\right) \\ &\times\left(U_j^h(B) - U_j^h(R)\right) \end{aligned} \tag{50}$$

Numerical optimization of η cannot actually use the simple product, because the product of two negative numbers is positive, so an optimization routine can achieve a large positive product by finding a position which is very bad for both actors. Therefore we use the simple product only in the quadrant where both terms are positive, employing a simple piecewise continuation over the other three quadrants that preserves continuity of values and slopes.

The Nash Bargaining Solution (NBS; Nash, 1950) has been criticized on the grounds that it seems to be a static model, without the dynamic bargaining so often observed. A similar situation is the contrast between von Neumann's Normal Form for representing a game and his Extensive Form (von Neumann and Morgenstern, 1944). The former is a matrix of strategic options, while the latter is a game tree of repeated moves and counter-moves. However, they are in fact mathematically equivalent. Similarly, it has been repeatedly shown that variations of the Rubenstein (1982) model of repeated rounds of non co-operative bargaining between agents leads to the same NBS, given suitable U and R.

In negotiation theory, the reference outcome R is often termed the 'Best Alternative to a Negotiated Outcome', or BATNA. In our case, the alternative (or BATNA) to a negotiated bargain between i and j is the conflict $i{:}j$, the utility of which we have already analyzed.

However, it is important to draw the distinction between i's estimate of the utilities, and j's estimates.

Recall that i selected j so i expected to gain by the conflict and expected j to lose from the conflict. In symbols, this means the following inequality will always hold, because i deliberately chose j to maximize the difference:

$$U_i^i(i \to j\ |\mathrm{S}) > U_i^i(S) \tag{51}$$

This almost always implies the converse inequality for i's estimate of j's utility:

$$U_j^i(i \to j\ |\mathrm{S}) < U_j^i(S) \tag{52}$$

Based on i's estimate of these two utilities, i might form an estimate of the best bargain it could get:

$$B^i = \arg\max_B \eta_{ij}^i(B|R) \tag{53}$$

KTAB Architecture: Formulae vs Submodels

KTAB itself has only top level formulae (such as adding up the votes in a coalition), with which the appropriately parameterized submodels are invoked. As just one consequence, there is no complete explicit formula in KTAB for the contribution of third parties, because parameterization enables dozens of submodels. Similarly, there is no formula in KTAB for the risk attitude of an actor, because it depends on the linear or square probability model, the naïve or strategic assessment, the range of R's employed, the adjustment of R and so on. Similarly, there is no formula for the probability of an actor succeeding or of a third-party being on the winning side. Some of the parameterizations of some of the submodels have closed form descriptions but most do not and must be numerically calculated. Hence, the simplest and least error-prone approach is to write high level code to numerically calculate everything, with appropriate parameterization for the various submodels as specified above.

As mentioned earlier, KTAB is a toolkit for building models. We have discussed how parameterization can be used to construct hundreds of different SMP models, and how this leads to an architecture in which only the top level abstract formulae appear explicitly. This architecture is further necessitated by the desire to support non SMP models of bargaining, for example the negotiations among managers over how to match individual staff members with particular projects.

Note that i's analysis takes into account its estimate of its own and j's utility, but does not represent making unnecessary concessions; it merely represents an assessment of the most i could get in hard negotiation with a rational opponent, given the intrinsic bargaining advantage which i has by virtue of expecting to win a confrontation. That is to say, i knows he has a favorable BATNA and expects j to have a similar belief.

In the SMP, a bargain is taken to mean a pair of new positions, one for actor i and one for actor j. While it may work out that these are the same, we do not assume that they always will be.

To use a basic example, such a negotiation is represented by an armed robber's demand for a victim's wallet. The BATNA is that the victim loses both wallet and life, while the robber gets both wallet and a murder charge. Thus, it is usually the case that both sides are better off than the BATNA if the victim surrenders his wallet-thus improving his outcome by staying alive-and the robber refrains from shooting, thus improving his outcome by avoiding a murder charge.

However, these are only i's estimates, which might or might not agree with j's estimates:

1. j might expect the inequalities to be reversed, so that j thinks it has the upper hand. This can lead to conflict;

2. j might agree that i has the upper hand, but think that the advantage is less than i does, so that j wishes to give fewer concessions than i expects;

3. j might agree that i has the upper hand, and think that the advantage is more than i does.

All these cases are subsumed in the problem of finding the NBS. The bargain expected is B_{ij}:

$$B_{ij} = \arg\max_{B} \eta_{ij}(B) \quad (54)$$

where

$$\eta_{ij}(B) = (U_i(B) - U_i(R)) \text{ x } (U_j(B) - U_j(R)) \quad (55)$$

If both actors are risk averse to any degree, there will usually be a solution because a deterministic bargain enables both sides to collect their risk premiums, even if they have to make small concessions to do so.

Note that the NBS is calculated using each actor's perspective on its own utility, so this represents not one actor's personal estimate of what the other would do but the outcome of a two-sided negotiation between self interested actors. A modeler could choose to simulate the step-by-step two-sided negotiation, similar to the Rubenstein approach, but for efficiency we calculate the NBS directly.

To take the third case first, suppose i expects to gain 10 points and j expects to lose 20 in the BATNA: a conflict will emerge. In this case, the NBS would be roughly that j gives 15 points to i, so that both sides are better off by their own estimate than they would have been in the BATNA.

In the first and second cases, the existence and detailed structure of the NBS depend on complicated interactions between the saliences along different dimensions. It may happen that i expects to gain 0.2 points of utility while j expects to lose 0.1. In one dimension with risk neutral actors, no NBS exists because it is not possible for i to gain 0.2 or more from j while j gives up 0.1 or less: j cannot move at least 0.2 toward i while simultaneously moving at most 0.1 toward i. However, if both actors are risk averse, then even in one dimension there may be enough of a risk premium to collect on both sides that a bargain exists. Similarly, in multiple dimensions, it may be possible for the actors to trade off gains and loses so that j gives to i what i values and j does not, and vice versa. This is similar to the situation of dividing vegetables and meat between a pure vegetarian and a pure carnivore: because they want different things, the goods can be divided up so that each gets everything they want.

KTAB offers two ways to do the bargaining.

1. The first is to literally carry out a numerical optimization in $2K$ dimensions to find the NBS. By using this method on thousands of problems, several patterns have been observed:

    - When neither actor is risk seeking (both have $0 \leq R$), the NBS exists and is always symmetrical in that both actors adopt the same position.
    - When one or both actors are risk seeking, an NBS may or may exist, and it may or may not be symmetrical. The more risk seeking the least risk tolerant actor is, the less likely it is that an NBS will exist. In other words, an NBS will usually exist even when one actor is highly risk seeking, so long as the other is fairly risk averse.
    - When both actors are highly risk seeking, and have very similar salience vectors, an NBS is unlikely to exist.
    - When neither actor is risk seeking, the NBS not only exists but is very closely approximated by a simple closed form formula. Remember that risk seeking behaviors are not expected among rational actors.

2. The second method is to use a closed form formula for a weighted average along each dimension k of the bargain B. Notice that we use the formula $P_i = P(i \succ j)$ with no superscript h, and similarly for P_j. This indicates that i's exertion of influence is calculated from i's perspective, while j's exertion is calculated from j's perspective, as the final bargain emerges from the interaction of two actors, each bargaining from their own perspective.

$$B_k = \frac{(s_{ik}P_i)^2\theta_{ik} + (s_{jk}P_j)^2\theta_{jk}}{(s_{ik}P_i)^2 + (s_{jk}P_j)^2} \tag{56}$$

As mentioned earlier, risk seeking behavior is never observed for significant stakes, so the second method is used in the example SMP application (with $0 \leq R \leq 1$).

Notice that while the R values do not explicitly appear in the interpolation formula, they are folded into the P values. The willingness of actors to compromise is affected by their R values: the cost to the Blue actor of a 0.1 shift in Figure 4.1 is one-tenth the cost to the Green actor, hence he will exert roughly one-tenth as much effort to resist it. Such a large difference in influence can shift the P values, particularly when every principal and third party have their own R values, and when they are all only partially adjusted from the principal actors' R values.

If the interpolation formula for B gives a positive value for $\eta(B)$, then the actors would find B mutually preferable to conflict. This results in sending 'bargaining proposals': i's proposal to j is that each adopt the same position, B.

If $\eta(B)$ is negative, then the actors have no bargain that is mutually preferable to conflict. This results in sending 'conflict proposals': i's proposal to j is that both adopt i's position, while j's proposal to i is that both adopt j's position.

4.4 Resolution

As mentioned earlier, every actor ends up with at least one and usually several bargains. Only one can actually occur; which one does is determined by the balance of influence, which in turn depends on the capabilities, utilities, and saliences of actors, i.e. on a PCE.

In this use of the PCE, the process being simulated is not one actor's analysis of future possibilities but the outcome when the various actors actually do exert influence to promote or oppose different outcomes. When doing so, each of the modeled actors would use its own perspective in voting for or against each of the various options.

Again, the PCE for resolution can be parameterized in several ways. First, the actors can vote according to any of the five voting rules. Second, they can assess the state that would result from each potential bargain by uniform weights or by an embedded PCE, i.e. by naïve or strategic voting. For the demonstration **smpc**, we used the proportional voting rule with naïve voting.

4.5 Simulation Turns

One step of bilateral negotiation yields a new position for each actor, which defines the next state. The entire process can be repeated until some termination criterion is met, yielding a full simulation. Two stopping criteria are currently offered:

1. A fixed number of turns; or

2. Until all changes are small compared with those which occurred in the first turn.

When choosing between the stopping criteria, the analyst can pick any simulation length, with any number of turns. Despite the formal framework we have presented here, mapping turns to real world

CDMPs is not straightforward. The exact length of time a turn represents is an abstraction. As suggested by the framework, a turn is any period of time during which all actors can exchange information and attempt to influence each other, through the four steps of assessment, targeting, proposal and resolution. There is a complicated and unresolved debate on how many turns are appropriate for a model like the SMP to run. Analysts will likely choose a number of iterations that balances the two major concerns: too short a simulation risks missing the value of an iterative CDMP model, while too long a simulation risks extending the results beyond what is credible based on a single dataset. By way of example, in Chinese SOE Reform we selected ten turns.

The final output from a run of **smpc** is a CSV record of the positions of each actor, as well as the probabilities for that state.

References

Black, Duncan. "On the Rationale of Group Decision-making," *Journal of Political Economy 56* (1948): 23-34.

Lester, Leo, Efird, Brian and Wise, Ben. “Reforming the Role of State-Owned Enterprise in China’s Energy Sector: An Analysis of Collective Decision Making Processes Using the KAPSARC Toolkit for Behavioral Analysis (KTAB).” *KAPSARC* (2015).

Nash, John. "Two-Person Cooperative Games." *Econometrica* 21 (1953): 128-140.

Nash, John. "The Bargaining Problem." *Econometrica* 18 (1950): 155-162.

Rubinstein, Ariel. "Perfect Equilibrium in a Bargaining Model." *Econometrica* 50 (1982): 97-109.

von Neumann, John, and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.

Wise, Ben, Lester, Leo, Efird, Brian. “An introduction to the KAPSARC Toolkit for Behavioral Analysis (KTAB) using one-dimensional spatial models” *KAPSARC* (2015).

About the Authors

Brian Efird

PhD, Senior Research Fellow and Program Director for Human Geography at KAPSARC leading teams on China, India, Local Content, Policy, and Bargaining Models.

Leo Lester

Research Fellow leading the North East Asia Research. Previously working in strategy and portfolio development for an international oil company, he has a PhD and is a CFA and FRM.

Imtenan Al-Mubarak

Senior Research Analyst focusing on collective decision making processes, energy, and policy research. She holds a Master's degree in Software Engineering from DePaul University in Chicago..

Ben Wise

PhD and a Senior Research Fellow at KAPSARC, working on models of collective decision making in the Policy and Decision Sciences Program.

Lama Yaseen

Research Analyst focusing on Energy Politics. She holds a BSc in Computer Science from Effat University in Jeddah, Saudi Arabia.

About the Project

KAPSARC is developing the KAPSARC Toolkit for Behavioral Analysis (KTAB), an open source software platform, to support modeling and analysis of collective decision making processes (CDMPs). Within our research, KTAB is intended to be the standard platform for analyzing bargaining problems, generalized voting models and policy decision making. It is our intent to use KTAB to assemble the building blocks for a broad class of CDMPs. Typical models in KTAB will draw on the insights of subject matter experts regarding decision makers and influencers in a methodical, consistent manner and will then assist researchers to identify feasible outcomes that are the result of CDMPs.

Notes

Notes

Notes

Notes

KAPSARC

مركز الملك عبدالله للدراسات والبحوث البترولية

King Abdullah Petroleum Studies and Research Center

www.kapsarc.org