Toward the Integration of Policymaking Models and Economic Models

Ben Wise, Brian Efird
King Abdullah Petroleum Studies and Research Center, Riyadh, Saudi Arabia
{ben.wise, brian.efird} @ kapsarc.org

Practical economic policy is often shaped by non-economic considerations. However, economic models typically treat policy choices as exogenously specified. Similarly, strategic models of bargaining typically use simple, fixed models of economics, if they use any economic model at all. In reality, political and economic policies are entwined as each may affect the other. This paper outlines one approach to endogenizing some of the non-economic interactions which shape the choice and implementation of economic policies. Economic policies can have disparate impacts across different regions, industries, or social groups. While the criterion of Pareto-optimality can direct policymakers' attention to the efficient frontier, there still remains the problem of weighing those disparate impacts to select a particular policy on the frontier. Understanding how different policymakers weigh the impacts differently is often key to estimating the range of plausible policy choices.

The KAPSARC Toolkit for Behavioral Analysis (KTAB) is a platform that enables the modeling and analysis of collective decision making processes (CDMPs), such as the negotiations over different economic policies. Rather than present the full range of models which can be built from KTAB libraries, we focus this discussion on a series of simple models, beginning with the classic one-dimensional spatial model of politics (SMP). The SMP is then generalized to handle multiple dimensions with different sub-models of actor's behavior. Next, the model is extended to include non-deterministic outcomes. With this framework in place, we then describe an early trial application of KTAB which links bargaining with a simple economic model, thus providing two initial examples of integrating the policymaking process and the economic consequences of policy choices. We build a very simple economic model to illustrate how natural search processes can generate strategically sophisticated economic policies.

This simple economic "model" is designed not to represent the best model of economic reality but to represent the expectations of actors; it is not an innovative result in its own right. It treats actors' policy position as a vector of taxes on *M* export goods. Subsidies are negative taxes, all goods are subject to tax or subsidy, and the total taxes must exactly match the total subsidies, given the constant-elasticity response of export demand to prices. Each actor represents either an industrial sector or a factor of production. The objective function of each actor is determined by the share of GDP they expect after imposing taxes. From the same export demand, sectors and factors have different economic expectations. The sectors have an economic view which includes projections of future steady state growth, so the effects on investment are accounted for, while the factors have an economic view which ignores investment effects. Both take into account distributional effects across sectors and factors, though they not only expect different consequences but also weigh them differently.

Keywords: Bargaining model, Domestic policy analysis, The GTAP Data Base and extensions, Collective decision making problems (CDMP), KAPSARC Toolkit for Behavioral Analysis (KTAB)

1 INTRODUCTION

Our proposed framework for integration is represented in Figure 1. This paper is organized in five sections, one for each major block of the diagram, a review of the integrated process, and finally conclusions.

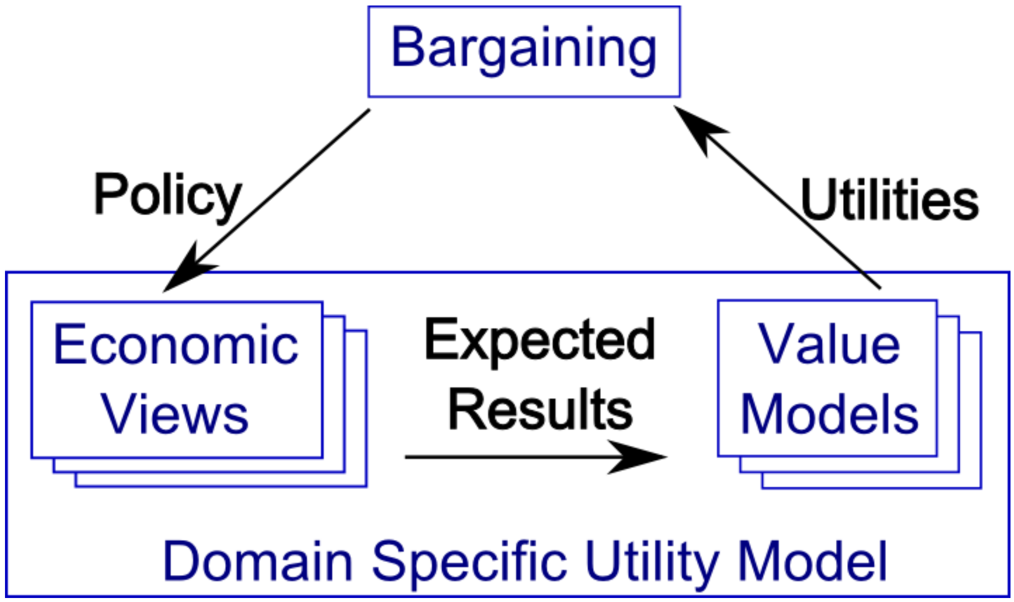

Figure 1: Integration Framework

The first and largest section is devoted to developing the bargaining model. Because the framework is intended to be usable with empirically based economic models, the bargaining model cannot assume a convenient mathematical form to ensure a tractable, closed-form solution. Therefore, we will develop a procedure which works for a broad variety of models. The fundamental integration is precisely to make the bargaining model flexible enough that a wide range of economic models fit in as subroutines. The motivations of the framework will be explained in stages, as we review some well-known, archetypal models for policy analysis. The overlaps with public choice theory and mechanism design will be discussed. At each stage, we explain the strengths of the archetypal model, as well as the limitations which this framework is designed to ameliorate.

The second section is a quick review of a simple Leontief model of an export driven economy. It is intended to model the estimated consequences which actors expect (rightly or wrongly) when debating policy options *prior* to implementation; it is not intended as the best available model of the actual consequences *after* implementing a policy. The example policy issue is to choose a vector of taxes and subsidies on exports. The economic stakeholders in the policy debate actually have different economic views and thus do not completely agree as to what results will be produced by a given policy. Some actors might derive their economic expectations from sophisticated, validated, state of the art economic models; some actors might use simple summaries of economic models; some actors might use only one criterion with no consideration of complex and indirect effects; some actors might be just plain wrong about the consequences of their favorite policy options. We take the economic viewpoints, or mental models, as given for better or worse, and examine how they interact in forming policy.

The third section describes how these particular actors assess the value to them of a policy. Because they occupy different positions in the economy, each actor would assign a different value to the same outcome (recall that not all actors will even agree on the outcome of a policy). In particular, we do not assume that there is an agreed-upon social utility function which is to be maximized.

2 Bargaining Models

Policy-making is often represented as the collective decision-making of a group of self-interested actors. To avoid detailed legacy interpretation of terms like “bargaining” or “negotiation”, we will generally term these “collective decision-making process” or CDMP. The CDMP stand in contrast to traditional economic models, which usually assume a large number of actors using a price mechanism in voluntary interactions.

One of the key characteristics is that there is a comparatively small number of policy actors, though they may judge their self-interest by the expected reactions of a large number of constituents, allies, opponents, etc. Like many bargaining models, KTAB is designed to work with a fairly small number (typically dozens, not hundreds, thousands, or more) of coarsely defined actors. An actor is usually not an individual person but a coherent group with common interests, such as “subsistence farmers” or “aerospace manufacturers”; individual leaders are taken as representatives of the groups they lead. This is entirely analogous to the use of a “representative household” or “representative consumer” in CGE models. Another key characteristic is that there is usually no price mechanism to mediate interactions (those interactions are handled by the economic model). While there are no prices, actors in CDMP are generally endowed with different levels of influence; examples include voting blocs of differing size, interest groups with more or less funds to support PR campaigns, shareholders with different numbers of shares in a corporate meeting, and so on. Actors can combine their influence to form coalitions, and a strong coalition can force an outcome over the objections of a weaker coalition; the classic example is an election where the candidate with the majority of votes wins, even though the minority did support and continues to prefer a different candidate.

Most models of policy making focus on the details of bargaining, and employ very simplified models of how actors derive utility from policies, so as to get analytically tractable solutions. KTAB takes the contrary approach of providing a toolkit to build simulation models whose solutions do not have simple analytical descriptions. Therefore, it is necessary to generalize the usual bargaining frameworks so that they still cover the classic cases, while also becoming usable for CDMP with less artificial models of utility – such as realistically complex economic models. One key tool for doing so is the concept of a probabilistic Condorcet election (PCE). Another is the concept of Central Position, which is useful for analyzing the convergence of a PCE over plausible economic policies.

2.1 Basic Notation

We briefly introduce some fundamental concepts; later sections will revisit them in more detail from several perspectives. There are N actors, each of which advocates some option θ_i. The set of possible options is denoted by Ω, so every option has $\theta \in \Omega$. The options which are near to θ in some sense is the neighborhood $\eta(\theta) \subseteq \Omega$. The set of positions currently advocated by the N different actors is the state, where $S \in \Omega^{\mathrm{N}}$. When we wish to emphasize which turn of the CDMP is being discussed, we add the time indices such as S_t.

In a choice between any two options x and y, actors can exert influence to try to get the group to adopt one or the other. The influence exerted is written as $v_i(x{:}\, y)$ where positive values favor x and negative values favor y. The simplest example is "one person, one vote" (1P1V): each actor exerts influence by casting a single indivisible vote for x $(v = +1)$, for y $(v = -1)$, or for some other option $(v = 0)$. Note that because the generalized vote is a net comparison of how much influence is exerted, $v_i(x{:}\, y) + v_i(y{:}\, x) = 0$.

Unlike personal utility, influence can be added. In the 1P1V case, this is simply adding up votes. The group "preference" for x over y is the following sum:

$$V(x{:}\, y) = \sum_{i=1}^{N} v_i(x{:}\, y)$$

Equation 1: Whole group vote from individual votes

We can split this into the total support for x and for y by looking at the influence exerted by supporters of either alternative.

The coalition supporting x over y is just the set of actors exerting positive net influence:

$$c(x{:}\, y) = \{i \mid v_i(x{:}\, y) > 0\}$$

Equation 2: Members of a coalition

We can write the strength of the coalition as the sum of the individual influences exerted to support x:

$$s(x{:}\, y) = \sum_{i \in c(x:y)} v_i(x{:}\, y)$$

Equation 3: Strength of a coalition

For the coalition supporting y over x, symmetric definitions of $c(y{:}\, x)$ and $s(y{:}\, x)$ apply.

In the case of 1P1V, $s(x{:}\, y)$ is just the sum of the votes supporting x, while $s(y{:}\, x)$ is just the sum of the votes supporting y, so $V(x{:}\, y)$ is just the excess of x's votes over y's. For more general cases, the option with the greater influence dominates an option backed by less influence:

$$V(x{:}\, y) = s(x{:}\, y) - s(y{:}\, x)$$

Equation 4: Whole group vote from coalitions

The 1P1V winner is the option x which has the most votes, i.e. $V(x{:}\, y) > 0$ for all other y. This can also be written as $x \succ y$ when we wish to emphasize that x is chosen over y, or as $y \rightarrow x$ when we wish to emphasize that x could follow y in a step-wise CDMP. An expression like $a \rightarrow b \rightarrow c \rightarrow d$ is a 4-element dominance sequence, where each option is dominated by the following one. It is well-known that majority rule frequently gives sequences with cycles, and that the larger the set of options S the more likely the sequence is to visit almost every option.

Traditionally, the option which dominates all others in pairwise contests is called the Condorcet winner; the entire process of determining $v_i(x{:}\,y)$, $s(x{:}\,y)$, and $V(x{:}\,y)$ for a particular state is sometimes termed a Condorcet election over that set of options. When the winner is determined strictly by $V(x{:}\,y) > 0$ regardless of how small is the margin of victory, we refer to the winning option as the deterministic Condorcet winner (DCW). As mentioned, it takes special conditions to avoid cycling so that a DCW exists.

An option X is the DCW if and only if, in every pairwise contest, the coalition favoring X is stronger than that favoring the alternative. Thus, a comparatively weak actor with many strong allies could easily prevail against a comparatively strong but isolated actor. The membership and strength of a coalition is affected by both intrinsic capability and the perceived stakes for each actor, which in turn are combined by the voting rule to estimate exercised influence.

In formal voting, even a tiny margin guarantees a victory. That is, a vote of 11:10 in a legislature means $V(x{:}\,y) = 11 - 10 = +1$ so first option wins, albeit by only one vote. However, when we consider informal voting as the exertion of informal influence, the outcome is better represented as probabilistic, where the probability of either option prevailing depends on the ratio of the coalition strengths. This leads to the concept of the probabilistic Condorcet Election (PCE), as we now discuss. We can define a probability that x is chosen over y simply by comparing the strength of coalitions:

$$P[y \to x\ \ ] = \frac{s(x{:}\,y)}{s(x{:}\,y) + s(y{:}\,x)}$$

This leads directly to a Markov process over the elements of the state S. Because the limiting distribution is determined by the whole state S, it is written as $P[x|S]$. In this paper, a PCE is the entire process of determining $v_i(x{:}\,y)$, $s(x{:}\,y)$, $P[y \to x\ ]$, and $P[x|S]$ for a particular state. We define the probabilistic Condorcet winner (PCW) as the option with highest probability from the PCE:

$$\arg\max_{x \in S} P[x|S]$$

For many common cases, the DCW and the PCW are the same option, but not always. Examples are discussed below.

We will now describe two common models of policy formation, the one-dimensional spatial model of politics and the Baron-Ferejohn model of bargaining in legislatures. We present them in terms of the framework just laid out above, so as to clarify both areas of similarity and of dissimilarity.

2.2 One-Dimensional Spatial Model of Politics

As described in (Black, 1948), the basic spatial model of politics (SMP) has three components:

- The set of plausible policy actions can be arranged on a line segment, where the position on the line specifies how much of something will be done. For mathematical convenience, this is generally represented as the unit interval, so $\Omega = [0,1]$.

- Interested actors exercise influence to try to produce a favorable policy action. The amount of influence they actually exert to promote one option over another depends on both their maximum capability to exert influence and the stakes for them in any given situation. This is their generalized vote, $v_i(x{:}\,y)$, where a positive value favors x and a negative one favors y.
- The preponderance of influence determines which policy will actually be implemented, as per Equation 1 and Equation 4.

In the original formulation, “exerting influence” was simply casting a formal ballot, and the “preponderance of influence” simply meant that the majority vote won. Black considered actors with differing voting-weights, w_i , similar to share-weighted voting in a corporate shareholder meeting. The 1P1V system sets all $w_i = 1$.

Thousands of political/military/economic situations have been successfully analyzed using the SMP [Feder 2002]; some examples include the following:

- Degree of government control of the economy, from complete laissez-faire to direct detailed control.
- Level of government surplus or deficit
- Whether foreign policy will be aggressively confrontational, passively accommodating, or something between.
- Whether a particular country will be tightly integrated with a multilateral military treaty (NATO), largely outside, or some intermediate level of engagement.
- Total support for one policy (or party), total support for the opposite policy (or party), or some level of compromise in between.

Detailed analyses of the entire voting history of the United States Congress using the SMP can be found in [Poole and Rosenthal, 1985, 1991, 1999, 2000]. The SMP was found to explain over 90% of the voting behavior.

An essential feature of the SMP is that the actors negotiate with each other, either explicitly or implicitly, formally or informally, to arrive at an outcome. There are several ways to implement dynamic bargaining in the SMP; a highly parameterizable SMP model called simply `smpApp` is included in the KTAB distribution along with several smaller examples of how to use the toolkit to build SMP models.

2.2.1 The Median Voter Theorem

One of the most famous and widely used results in voting theory and analytical politics is the Median Voter Theorem (MVT), originally published in (Black, 1948). We discuss it to illustrate what features are generalized to handle our Leontief example, and what features are retained.

Black outlined a series of assumptions under which a specific negotiation outcome could be easily computed. The first set of assumptions specify the structure of the SMP; the second set fill out the CDMP.

Every actor prefers alternatives near their position over alternatives further away. It is critical for the MVT that alternatives can be ordered by distance and that all actors agree on the order. Further, the intensity of the actor’s preference between two alternatives increases as the distance between them increases. In other words, $u_i\ (\theta)$ is single-peaked at θ_i. For convenience, the utility of options is typically normalized to the von Neuman scale of [0,1].

When faced with a choice between alternatives, each actor will exert all their influence to promote the one they prefer, no matter how small the difference between the alternatives. This "all-or-nothing" response is termed "binary voting". This behavior is described by the following equation:

$$v_i(x{:}\,y) = \begin{matrix} +w_i, & \text{if } u_i\ (x) > u_i(y) \\ -w_i, & \text{if } u_i(x) < u_i(y) \end{matrix}$$

Equation 5: Binary voting

These assumptions are enough to determine the DCW, which is called the weighted median position (WMP), which has two properties: at least half the actors' total influence is at or to the left of the WMP, and at the same time at least half is at or to the right. The key feature of the WMP is that no coalition would either be able to or be motivated to influence the group to choose some other alternative over the WMP.

While the WMP does have the property of stability, given the balance of power and interest, it is not obvious that the WMP will ever enter the CDMP, unless one of the actors just happens to have it as their ideal position initially. Therefore, Black introduced a second set of assumptions to characterize the process of dynamic bargaining over alternatives:

- The CDMP takes discrete steps, as proposals are made to the group sequentially.
- After each proposal, the group votes between two options: whether to accept the new proposal or to keep the old one. Whichever wins is accepted, temporarily.
- Regardless of the outcome of a vote, the bargaining process continues and some actor makes a new proposal.
- The bargaining process stops when no actor finds it advantageous to make a proposal, at which point the last proposal accepted is the one to be implemented.
- Actors always vote between alternatives based on which they prefer most, regardless of who proposed them or why.
- Some actors propose alternatives based only on how likely they are to be selected by the group.

While this process was presented as an explicit model of formal committee voting procedures, it has been used very successfully to modeling the implicit bargaining of more informal processes, such as political campaigns, back-room persuasion before parliamentary voting, price-setting within an oligopoly, civil unrest, and so on. The term "generalized voting" is used to cover both formal and informal exertion of influence.

Black's MVT rests on the observation that, if it were ever proposed, the WMP would dominate any other option in a pairwise choice, so it is the DCW. This situation is illustrated in Figure 2. There are seven actors, with total strength 97. Actor D (strength 15) has total 53 at its position or to the left, while simultaneously having 59 at or to the right

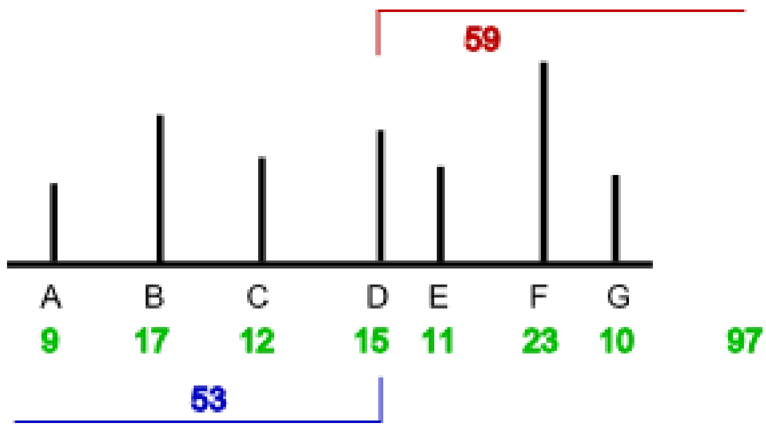

Figure 2: Weighted Median Position

Suppose the whole group faced the choice D:E. All those at or to the left of D would prefer D to E, and they exert their maximum influence, so $s(D{:}E) = 53 = 9 + 17 + 12 + 15$. All those at or to the right of E would prefer E to D, so $s(E{:}D) = 44 = 11 + 23 + 10$. So $V(D{:}E) = 53 - 44 = +11$. Similarly for all the other advocated positions.

It is crucial to this argument that all actors agree on the left-to-right ordering of the options, and that their vote depends only on the ordering, not the precise distances. Thus, Ω must be one dimensional and actors must use binary voting.

Once the WMP it is proposed, all the actors know that it is the DCW and no proposal could displace it, so further proposals are pointless, and the process stops.

Note that because the WMP will be implemented after it wins the final vote, whoever proposed the WMP is the one whose policy will be implemented. Further, they can vary their proposal within narrow bounds and thus adjust the final policy as they wish. Thus, the weighted median actor can be seen as decision-maker. However, they can only adjust the final policy within certain bounds, subject to the provision that they don't alter the policy so much as to lose support of their coalition, i.e. that it remains the WMP. Again, this is illustrated in Figure 2. As long as it stays between actors C and E (strength 12 and 11), actor D remains the WMP. But if it moves past either, then coalitions shift and the one it passed becomes the WMP.

Black considered the case of a few actors who craft their positions in light of the expected actions of a continuous distribution of voters. These actors are generically called "office-seeking politicians" (OSP), though they can represent a broader range of actors whose main goal is to be the decision-maker. Under these conditions, the WMP is that point with exactly half the voters on either side. Office-seeking politicians can do the preceding strategic analysis and foresee the result, so they will always have a motive to make a proposal closer to the WMP than was the most recent proposal. It is the competition between the OSP which ensures that the CDMP will move toward and finally reach the WMP, at which point it will stop.

The bargaining process can be formulated by looking at how the DCW of a set of positions changes over time. Black describes an iterative process where there is a proposal, z_t, "on the floor" at each step, and actors propose amendments to it. As explained, they will be motivated to propose amendments which dominate the current proposal, and the dominant amendment will be selected:

$$z_{t+1} = \underset{x \in \eta(z_t)}{\operatorname{argmax}} V(x: z_t)$$

Equation 6: Basic DCW bargaining process

The essence of the MVT is that under all the required assumptions, this iterative process converges to the WMP and terminates.

It is well-known that MVT does not generalize easily; some of the reasons are the following.

- In multiple dimensions, each dimension has a different weighted median, hence a different winning actor.
- When the utility function is not single-peaked at the actor's position, actors might find distant options more appealing than nearby ones, so it is impossible to tell how they will vote based on mere proximity.
- When options cannot be placed on a line at all (e.g. combinatorial matching problems), then the definition of the WMP become problematic, because one can no longer assume that all "nearby" options dominate all "far away" options.

Unfortunately, economic problems often involve all these complications, and more. In our simple Leontief example, we have multiple tax rates, so the actor's positions cannot be arrayed on a single dimension. Further, the utility of a tax vector to an actor depends not on the difference between the tax vectors but on the difference in value-added in their sector after those taxes are imposed. Thus, it is perfectly possible to alter all the components of a tax vector while keeping the value added for a particular actor constant. Small variations in each such vector – no matter how near or far from the actor's position - can increase or decrease their value added, which proves that their utility function is not single-peaked.

Because there is a simple solution for the final outcome of Black's model, it is frequently forgotten that the solution is based in a dynamic CDMP, as per Equation 6. This is important because the underlying CDMP can be adapted to other models in order to obtain useful results even when no closed form exists; this is one of the kinds of CDMP which KTAB is designed to simulate. The main advantage of the simulation approach is that we do not need to constrain problems simply to obtain a closed-form solution. The main disadvantage is that complex economics models (such as IO, SAM, and CGE), do not have the simplified and constrained structure necessary to apply the SMP, so the more general approach of Section 2.1 is required.

2.3 Generalized Voting and Voting Rules

We have already described binary voting in the context of the MVT. The KTAB toolkit offers several rules for generalized voting; model builders are free to define more. This flexibility has profound implications for the design of KTAB and kind of problems it can address, as the choice of voting rule completely changes the bargaining process.

The system of generalized pairwise voting is quite easily extended to any system where actors assign scores to options, such as range, approval, or Borda voting, simply by taking the difference between scores. Similarly, there are simple extensions to super-majority voting, highest two options win, and so on. These extensions are described in (Wise, Lester, and Efird, 2015).

In many situations, binary voting is too extreme and a more nuanced application of informal influence is common. Proportional voting is designed to model the situation where there is a wide choice of means of exerting influence, an implicit budget on exerting influence, an intrinsic cost to exerting influence, or all three. Typical examples include spending money to fund political relations campaigns, nations deploying military force to pressure other nations, special interest groups lobbying for tax and subsidy policies, and so on. Obviously, political parties and special interest groups expend their campaign funds carefully, expending little or no resources on minor issues and making the largest expenditures for the most important issues, so binary voting is not an accurate model of their behavior.

Proportional voting is the exertion of the actor's influence in proportion to their perception of the stakes, which is how much they stand to gain or lose from the two alternatives. The proportional voting rule is as follows:

$$v_i(x:y) = w_i\,[u_i(x) - u_i(y)]$$

Equation 7: Proportional Voting

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

Proportional voting is not only intuitively plausible, but it has very desirable mathematical properties which will be explored below.

Binary and proportional voting are not the only possibilities. There is a well-known political rule that policies should "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. The cubic voting is designed to model precisely this situation where actors respond little to small changes, but respond energetically both to large losses and to large gains:

$$v_i(x:y) = w_i\,[u_i(x) - u_i(y)]^3$$

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

One can moderate the nonlinearity of the binary and cubic voting rules by taking a weighted average with the proportional voting rule. The proportional voting rule is contrasted with two such hybrid rules, Binary-Proportional, and Cubic-Proportional, in Figure 3.

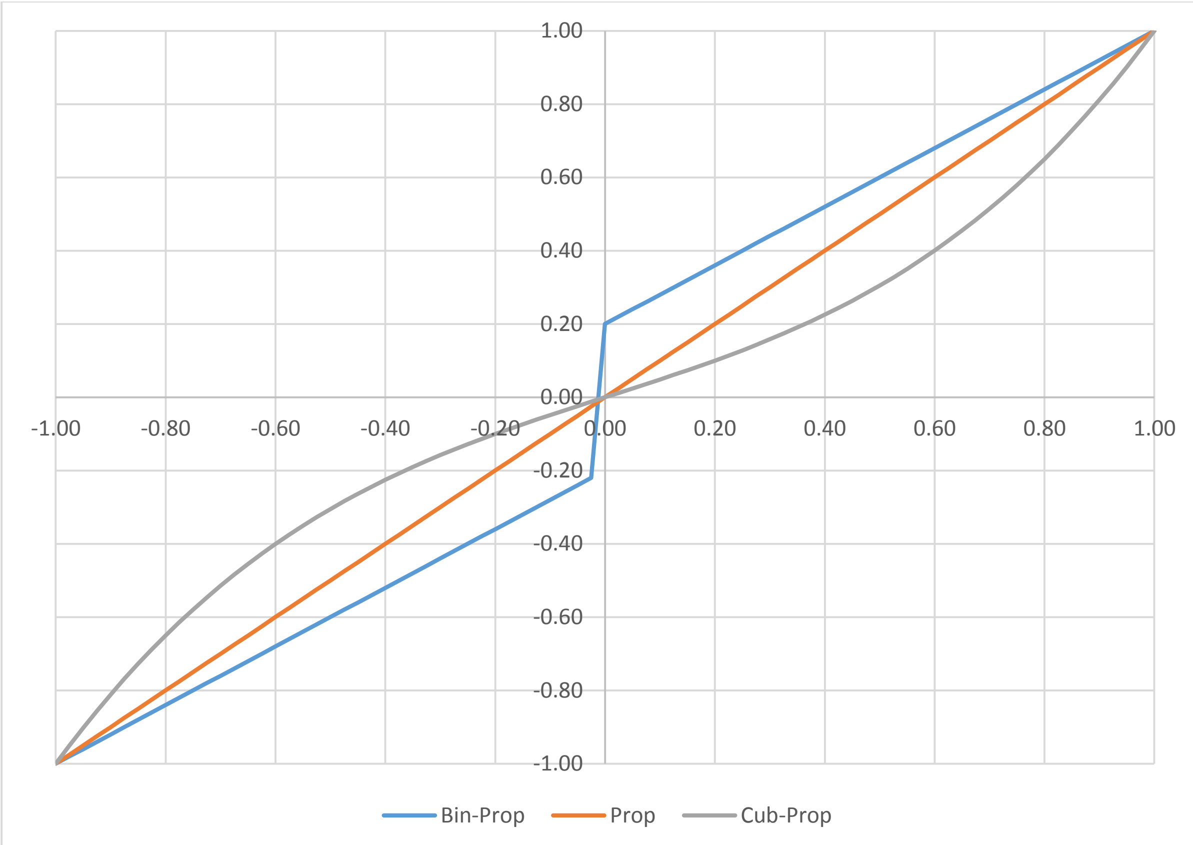

Figure 3: Binary-Proportional, Proportional, and Cubic-Proportional Voting Rules

The point of mentioning the cubic and hybrid rules is that while they have some empirical plausibility, they do not yield a DCW for any Ω or u known to these authors. However, as mentioned earlier, there is anecdotal evidence that some situations are better described by a proportional-cubic rule than binary or proportional, so a toolkit to model such situations cannot rely on the existence of a DCW. Therefore, the KTAB toolkit does not assume any particular voting rules; in particular, it does not compute either weighted medians or central positions. Instead, KTAB relies on the more general concept of a PCE. A useful result for understanding the behavior of PCE is the Central Position Theorem.

2.4 The Central Position

The Central Position Theorem (CPT) was derived in [Wise 2010a, 2010b] and published in [Jesse 2011]; a closely related but distinct case is analyzed in Corollary 4.4 of [Coughlin 1992]. The CPT states that under proportional voting, the CDMP will lead to an outcome which maximizes a function we call the Weighted Attractiveness Score (WAS). That outcome is called the Central Position (CP), and it is the outcome which the whole group's balance of influence will favor over any other outcome. In other words, it is the DCW under proportional voting. By almost the same logic as the MVT, office seeking politicians will always have an incentive to propose options with higher WAS, until the CP is proposed. At that point, no other option can displace the CP and no further negotiation will occur, because every actor knows that any further proposals will be rejected. This situation can be expected to persist until something changes the balance of Influence. In the SMP-1 with binary voting, the detailed shape of an actor's Attractiveness curve is irrelevant: actors exert their full voting weight in favor of whichever alternative is closer to their own Position. With proportional voting, the shape of the Attractiveness curve is necessary to compare the attractiveness of two options and thus gauge how much Influence an actor will exert for one or the other. The logic of the CPT is fundamentally non-spatial, which makes it very broadly useful but also harder to visualize. The strong assumptions which made the MVT easy to visualize are no longer used, so there is no simple geometric explanation of the CPT. This is intimately connected to the fact that the CPT applies not only to one-dimensional or multi-dimensional problems but also to discrete combinatorial problems in which notions of distance are not useful. However, the proof of the CPT is quite simple algebraically, as we shall see shortly. What is important for our purposes here is that it can easily be applied to economic CDMPs with multiple numeric parameters, rather than just the SMP-1.

Coughlin's corollary considers the situation where the option set Ω is continuous, convex, and compact; every voter casts weighted binary votes according to a binary Luce model; each voter has a utility function of a specific exponential form; and expected net votes are estimated by a probabilistic voting model. His corollary 4.4 states that the electoral equilibrium is that which maximizes an implicit Benthamite social utility function; for brevity, we call that position simply the Central Position. Note that corollary 4.4 initially uses a probabilistic voting model, but then assumes the actual votes are exactly the mean expected vote. This assumption can be seen as a way to bridge the gap between discrete voting subject to Arrow's Impossibility Theorem (Arrow, 1950) and continuous, proportional voting subject to the Central Position Theorem.

We now present the proof that, under proportional voting, a DCW exists; we define the CP to be that particular DCW. As mentioned earlier, the balance of the whole group is the sum of the individual influences exerted:

$$V(x:y) = \sum_i v_i(x:y)$$

Under proportional voting,

$$V(x:y) = \sum_i w_i \, [u_i(x) - u_i(y)]$$

We can factor the right hand side, because the sum of the differences is the difference of the sums:

$$V(x:y) = \sum_i w_i \, u_i(x) - \sum_i w_i \, u_i(y)$$

We can see that the CP is the option which maximizes the following WAS:

$$\omega(x) = \sum_i w_i \, u_i(x)$$

Equation 8: Weighted Attractiveness Sum

In [PC92], this WAS is termed the Benthamite social utility function. Notice that the WAS is not simply postulated as a function to be maximized but derived from an underlying model of a bargaining process. This is entirely analogous to the total surplus which is maximized in markets: maximization of total surplus is a logical consequence of a particular model of dynamic interactions, not an assumption.

The MVT states that the SMP-1 with binary voting gives the weighted median position as the DCW. In contrast, the CPT states that proportional voting alone gives the central position as the DCW, with no further assumptions about the shape of the utility function or the set of options. (Some technical assumptions about open sets and limits are needed for both the MVT and the CPT). Because it is impossible to have both $\omega(x) > \omega(y)$ and$\omega(y) > \omega(x)$, the "group preference" is transitive and acyclic with proportional voting.

2.5 Polarization versus Compromise

Because $\omega(x)$ can be multi-modal, the CP can jump discontinuously in response to small changes in actor's utility or capability, which does not happen with the WMP. While there are many analyses of why political parties do not always coalesce to the WMP in the one dimensional case, and why political "tipping points" can lead to discontinuous change, the CPT provides a particularly succinct explanation of both phenomena. As this point is relevant to our economic bargaining example, we will explain how it arises in the following simple one-dimensional case.

We consider six actors. The position of each actor is characterized by a number in the unit interval, so $\theta_i \in \Omega = [0,1]$. The positions might correspond to something like a left-right political spectrum, level of government debt, or some other one-dimensional policy issue. The following table gives the positions of the actors, as well as the w_i values for maximum possible influence.

Position	Influence
0.13	23
0.19	15
0.25	74
0.84	71
0.89	27
0.93	13

Table 1: Positions and Influences

Under the proportional voting model of Equation 7 and Equation 8, the central position (CP) is that which maximizes the weighted attractiveness sum. The behavior of the CP depends crucially on the shape of the actors' utility functions. While utility functions of goods are required to be increasing and concave, the utility functions here are not in terms of goods but in terms of political policies. Hence, there is no reason to expect them to be either concave or convex. We compare the cases of convex and concave, where the positions and influence of the actors are held constant at the values of Table 1.

If the utility functions of these actors are convex and single-peaked at the positions of Table 1, then we get the jagged, two-peaked shape in Figure 4. Clearly, any small changes in the curve can cause the CP to jump discontinuously between 0.84 and 0.25. Note that it is always a position current held by one of the actors. Due to the convexity, the group will select one or the other, without compromise.

Given that not all actors will have precisely the same estimates of policy parameters, they will have slightly differing estimates of the "altitude" of each peak. It is easy to draw a band of uncertainty crossing the two peaks and see that the OSP will separate into two distinct clusters, one around each peak.

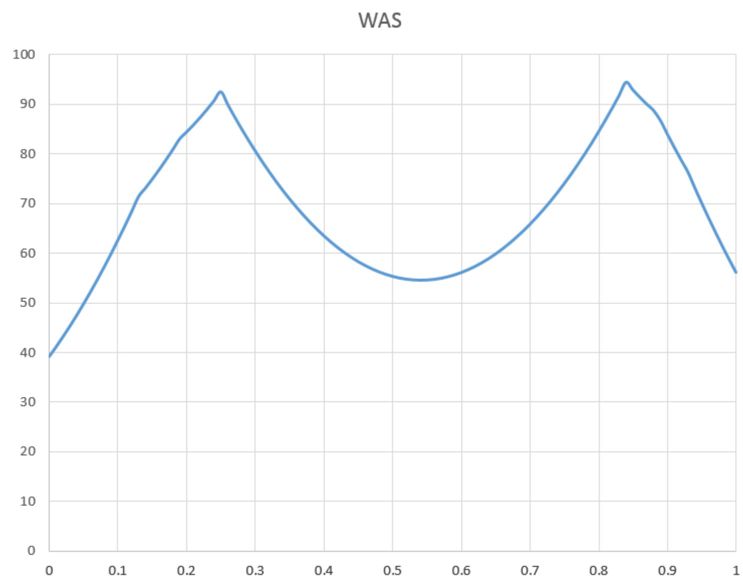

Figure 4: Weighted Attractiveness Sum with Convex Utility

On the other hand, if the utility functions are concave and single-peaked at the positions of Table 1, then we get the smooth, single-peaked curve in Figure 5. The CP happens to be at 0.54 for this data, a compromise position not currently held by any of the actors. Due to the concavity, the group will compromise and “split the difference”. Further, variations in the data cause the CP to shift smoothly as the single peak moves, rather than jumping discontinuously as in the prior case.

Again, because actors will have different estimates of political parameters, they will choose different positions of slightly differing “altitude”. In this case, the band of uncertainty will cover a single range around the CP, so the OSP will form one cluster around the CP.

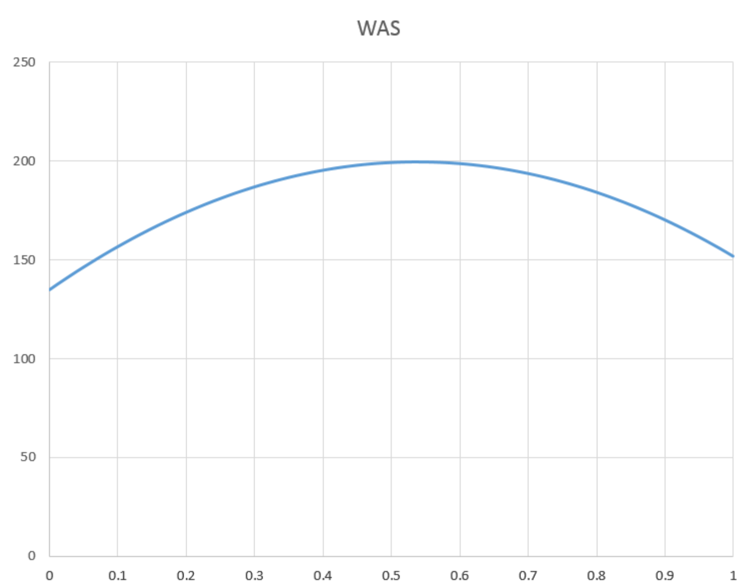

Figure 5: Weighted Attractiveness Sum with Concave Utility

In our Leontief economic example, the actors have utility functions which are basically concave in their particular tax, so we expect the CP to be more of a compromise, and the example does exhibit this behavior.

In one dimension, it is simple to locate the CP numerically by scanning through the values. However, in general the maximization problem to locate the CP can range from very easy to very difficult, because the CPT does not impose stringent conditions on the option set Ω or the utility functions u_i . Even when $\Omega \subseteq R^N$, the problem can be quite difficult if ω is multi-modal and Ω is nonconvex. When Ω is a set of discrete combinatorial structures without even a distance metric (for example, the decision of which regions will or will not participate in a carbon emission trading scheme (ETS), or which set of regulations will govern the ETS) the maximization problem is even more difficult. For this reason, KTAB includes some simple optimization routines, such as genetic optimization.

2.6 Baron-Ferejohn Model

The discussion so far has focused on spatial models, but there are some very well-known non-spatial models of bargaining. The Baron-Ferejohn Model (BFM) as presented in [Baron, Ferejohn 1989] addresses the problem of multiple parties dividing up a fixed resource. This may be used to model the formation of a government in a multiparty parliamentary system. It can also be used to model representatives of different regions dividing up the national budget. In the latter interpretation, it is a simple model of policy formation, with a trivial economic model.

Suppose there are N actors, each with a voting weight w_i. In the BFM, the position of an actor is a vector of fractions, indicating how much of the resource goes to each actor. Thus, with five actors, the position $[0.37, 0.23, 0, 0.40, 0]$ would be the proposal that actor 0 gets 37%, actor 1 gets 23%, actor 3 gets 40%, while actors 2 and 4 get nothing. In the SMP-1, the position is a single number, regardless of how many actors are involved:$\Omega = [0,1]$. In the BFM, a position is an N-dimensional vector. Because of the constraint that the fractions add to 100%, it is a non-negative N-dimensional vector with N-1 degrees of freedom:

$$\Omega = \left\{ (x_1, x_2, \ldots, x_N) \mid \forall i \; x_i \geq 0, \qquad 1 = \sum_{i=1}^{N} x_i \right\}$$

The BFM uses a CDMP somewhat similar to the one Black used to derive the MVT. The CDMP proceeds in a series of turns, analogous to a series of amendments proposed in a legislative session. At each turn, an actor is chosen to propose a position, i.e. a division of the resource. The whole group of actors then vote yes or no on the proposal, where each actor has a binary weighted vote. If the weighted majority favors the proposal, it is adopted and the CDMP terminates. Otherwise, an actor is chosen with uniform probability and the process repeats.

At each turn, each actor can assess the expected value to them of the other actors' positions (including whatever self-interested position they might adopt). Because the process might or might not continue after each vote, the assessment must look at the expected value of the game tree. The expected value of letting the process continue (given all the intelligent strategic behavior of all the other actors) is termed the continuation value. While there are more variations of the BFM than can be reviewed here, the essence of the BFM is to determine the minimum winning coalitions and the limiting, strategic equilibrium value of the division-vector. No actor will accept a proposal that gives them less than their continuation value, and none will offer a proposal which gives away more than their continuation value. Of course, the predicted outcome is the DCW for weighted binary voting over that set of positions, using the continuation values as the utility function. It can be computed by an iterative adjustment procedure (rather like computing an eigenvector). When the iterative adjustment procedure is implemented as a CDMP in KTAB, the same result is obtained.

The BFM does not include actors losing patience and leaving the negotiation. There is no reversion to a pre-defined status quo in case negotiations break down. The process repeats without limit until the actors accept a proposal. There is no penalty for taking many turns. The only actions available to actors are those specified in the mode: to vote on proposal, or to present a proposal. The model does not include actors resorting to threats and inducements not included in the model. There are no side payments. Much of the content of [BGM89] is a detailed algebraic proof that the process as defined in the paper does terminate promptly.

While there are many variations on the BFM (e.g. actors are selected with probability proportional to their voting weight), there are several characteristics we wish to emphasize.

In the BFM, the complications of real economic models have been suppressed in order to focus on the complications of negotiations. The trivial economic model is that policy choice directly specifies economic outcomes for each actor: if a position specifies a 0.23 share for an actor, then their utility is precisely 0.23. By contrast, the SMP uses an intermediate function to map positions into utilities. In the SMP-1, the utility declines with distance, and all actors measure distance along the line in the same way. If the SMP-1 actors were to sort one-dimensional positions by how close they were to a fixed point, all actors would agree on the order. In contrast, each actor in the SMP-N might have their own unique measure of weighted Euclidean distance, depending on how they weigh the different dimensions of the problem. Thus, if SMP-N actors were to sort multi-dimensional positions by how close they were to a fixed point, they might not even agree on the order. In this sense, the actors in SMP-N have very different views of how the space of options is structured.

These examples illustrate why KTAB is structured so that the user can specify a Domain Specific Utility Model (DSUM). In our notation, $u_i(\theta)$ is the value from the DSUM for the utility to actor I of position$\theta \in \Omega$. The SMP-1, SMP-N, and BFM all use fundamentally different models of how policy positions map into the utilities of actors, but all are quite simple compared to realistic economic models.

2.7 PROBABILISTIC CONDORCET ELECTION

As mentioned earlier, a committee vote of 11:10 between two options is a clear win for the first option. However, we also consider informal voting, which is more free-form exertion of influence. In cases of informal exertion of influence, such as back-room persuasion before parliamentary voting or expenditure of funds to support a political campaign, the outcome is better represented as probabilistic. If two rival political campaigns each spent ten million dollars to support their candidate, then *ceteris paribus* each candidate would have 50% chance of victory. If they could fund campaigns in the ratio

11:10, the first candidate would not be absolutely guaranteed victory, though the probability would be raised somewhat above 50% *ceteris paribus*.

A common approach to this situation is to model the victory probability as the ratio of the strengths of the coalition supporting each option. Thus, the probability that the group will choose x over y increases as the coalition supporting x over y strengthens:

$$P[y \to x] = \frac{s(x{:}\, y)}{s(y{:}\, x) + s(x{:}\, y)}$$

Similarly, the probability of the group choosing y over x increases with the strength of y's coalition:

$$P[x \to y] = \frac{s(y{:}\, x)}{s(y{:}\, x) + s(x{:}\, y)}$$

If these probabilities are interpreted as transition probabilities, we immediately get a Markov process. The steps of the hypothetical Markov process play the same role as the steps of the hypothetical bargaining process in the MVT or BFM: an analytical approximation to a much less orderly real-world process. The limiting distribution,P, of the Markov process gives the likelihood that each particular outcome is the selected option at any given turn of the CDMP. It can easily be obtained by the iterative procedure of starting with a flat distribution and multiplying it by the transition matrix until the distribution stabilizes. For the analytically tractable case of a state with just two options, $S = (x, y)$, the limiting distribution is exactly the intuitive result:

$$P[x|S] = \frac{s(x{:}\, y)}{s(y{:}\, x) + s(x{:}\, y)}$$

$$P[y|S] = \frac{s(y{:}\, x)}{s(y{:}\, x) + s(x{:}\, y)}$$

The entire process of computing the $s(y{:}\, x)$ matrix, the transition probabilities, and the limiting distribution, is termed a Probabilistic Condorcet Election. The option with the highest probability from the PCE is termed the Probabilistic Condorcet Winner (PCW). The PCW has important similarities to a standard DCW, as well as some useful differences.

For SMP-1 with binary voting, one can compute the coalition structure to determine both the DCW and the PCW. Almost always, the PCW will be the DCW, i.e. the WMP.

For a wide variety of problems (SMP-1, SMP-N, discrete combinatorial with random utilities, and more) with proportional voting, one can compute the coalition structure to determine both the DCW and the PCW. Almost always, the PCW will be the DCW, i.e. the CP.

Nevertheless, one can construct situations in which the PCW is not the DCW. The simplest examples have a few carefully-chosen net votes which are orders of magnitude below the others. Suppose there are four options, A through D. The strength of each coalition $s(x{:}\, y)$ appears in the x-th row and y-th column. The strength $s(A{:}\, B) = 10.01$ appears at the intersection of row A and column B; the strength $s(B{:}\, A) = 10.00$ appears at the intersection of the row B and column A; the net vote is $V(A{:}\, B) = 0.01$. Similarly, the net vote $V(B{:}\, C) = 1.0$. The net vote $V(A{:}\, B)$ is two orders of magnitude below $V(B{:}\, C)$.

$s(x:y)$	A	B	C	D
A	0	10.01	10.01	10.01
B	10.00	0	11.00	11.00
C	10.00	10.00	0	10.01
D	10.00	10.00	10.00	0

Figure 6: Coalition Strengths

As we intend to illustrate the effect of making one net vote become very small, we define $\epsilon = 10^{-2}$. The net votes can be derived from the s matrix as follows:

$$V(A:B) = s(A:B) - s(B:A) = (10 + \epsilon) - (10) = \epsilon$$

$$V(A:C) = \epsilon$$

$$V(A:D) = \epsilon$$

$$V(B:C) = s(B:C) - s(C:B) = 11 - 10 = +1$$

$$V(B:D) = +1$$

$$V(C:D) = \epsilon$$

The DCW is A, because it has a positive, albeit small, margin of victory over all others. B significantly defeats options C and D but slightly loses to A. If A's net vote against B, C, and D remains positive but shrinks to negligible values, e.g. $\epsilon \to 10^{-20}$, then A effectively ties B, C, and D yet remains the DCW. However, B is the PCW because it defeats C and D by significant margins (+1 each) and effectively ties with A. When$\epsilon = 0$, then B remains the PCW but suddenly becomes the DCW, as it loses to no option and defeats some.

The contrast between the DCW and the PCW arises from the fundamental difference in how the treat $V(x:y)$: the former depends discontinuously on only the sign, while the latter depends continuously on magnitudes. When comparing integer votes in committees, the DCW and PCW agree because miniscule percent differences in coalitions cannot occur: the coalition strengths are integers of a few hundred or less. However, when informal influence is described by real numbers, and even 64-bit floating point arithmetic introduces random errors on the order of10^{-20}, these pathological cases can arise.

The DCW and PCW are clearly not just different ways of computing the same result. However, they are similar in a common and important case: when the DCW does exist, and the margins of victory are all within a few orders of magnitude of each other, the DCW almost always coincides with the PCW.

Another fundamental difference is that the PCW always exists, even when the DCW does not. This is because, for any matrix of positive strengths, $s(x:y)$, the corresponding transition probabilities do exist and do define a Markov process with a limiting distribution, i.e. a PCE. Given a few technical conditions on open sets and limits, the distribution will have a maximum, i.e. a PCW. For example, the SMP-1 generally does not have a DCW under cubic voting, even though it always has a DCW under both binary voting (the WMP) and proportional voting (the CP). Empirically, the limiting distribution will be much flatter when there is a cycle of options than when there is a single DCW.

This observation touches on the third fundamental difference. Unlike the DCW, the PCE provides not only a point-estimate of the most likely outcome (i.e. the PCW), but also an indication of the dispersion of plausible results around that estimate. For example, applying a PCE to the SMP-1 with cubic voting generally yields a much flatter distribution than when it is applied to the SMP-1 with binary voting.

One way this result can appear is when a few options dominate all others, but the group has circular preferences between those few non-dominated options. Even in this case, the PCE can provide useful information beyond the simple statement that there are circular preferences. The limiting distribution assigns non-zero probability to each non-dominated outcome, and the varying probabilities among the non-dominated outcomes can provide useful information if they vary drastically.

2.8 EXPECTED VALUE OF A STATE

We define the state of a CDMP during a given turn as the set S of positions θ_i taken by the actors at that turn:

$$S = (\theta_1, \theta_2, \dots, \theta_N)$$

Given the DSUM on positions,$u_i(\theta)$, and a voting rule of each actor, we can build the strength of the coalition supporting each option in any given pair. As this is just another choice between two options, we construct coalitions and strengths exactly as before:

$$c(\theta_j : \theta_k) = \{i \mid v_i(\theta_j : \theta_k) > 0\}$$

$$s(\theta_j : \theta_k) = \sum_{i \in c(\theta_j : \theta_k)} v_i(\theta_j : \theta_k)$$

For the purposes of demonstrating integration of policy-making and economic models, we assumed proportional voting. The transition probabilities can then be estimated, so the limiting distribution P_∞ of the PCE can be obtained.

We define the value of an entire state as the expected utility to an actor of the outcome of the state:

$$U_i(S) = \sum_j P[\theta_j | S] \, u_i(\theta_j)$$

If actor k were to change their position to α, we denote the new state $S' = \sigma(k, \alpha | S)$ with just one position changed. We denote the expected utility of the new state as $U_i(S')$, where the PCE is repeated in the new state. These expressions can be combined, so the expected utility to actor I of the new state $is\ U_i(\sigma(k, \alpha | S))$, where the PCE is repeated in the new state.

This leads to CDMP definition where each actor separately crafts a new position to maximize the expected utility to themselves, given the support or opposition they expect from other actors:

$$\theta_{i,t+1} = \underset{\alpha \in \eta(\theta_{i,t})}{\operatorname{argmax}} U_i\big(\sigma(i,\alpha|S_t)\big)$$

Equation 9: Basic EUMax bargaining process

Supporting or opposing a policy based solely on the utility expected if it were to be implemented is called "naïve voting". With proportional voting, a naïve voter would exert influence in a choice between two positions as follows:

$$v_i(\theta_j \colon \theta_k) = w_i\left(u_i(\theta_j) - u_i(\theta_k)\right)$$

Thus, a naïve voter would exert the most effort to support an outcome they most prefer (compared to the alternative), and the least effort to support an outcome they did not greatly prefer.

Supporting or opposing a policy based on the expected outcome of the entire state is one type of "strategic voting". With proportional voting, a strategic voter would exert influence as follows:

$$v_i(j,\alpha : k,\beta) = w_i\left(U_i\big(\sigma(j,\alpha|S)\big) - U_i(\sigma(k,\beta|S)|S)\right)$$

In particular, the sign $v_i(j,\alpha : k,\beta)$ would indicate whether actor i found it more strategically desirable for actor j to adopt position α or for actor k to adopt β, given the current state S. Note that there are nested voting models: naïve voting is applied within a state, while strategic voting is applied across states.

Note that the CPT applies to strategic proportional voting over possible states because the CPT requires no special structure over the set of options or the utility functions. We simply define the set of options to be the (very large) set of possible states and the utility function to be expected utility from the PCE, and the CPT automatically applies. This means that, with proportional voting, there is a "central state" that the group will adopt over any other, and that the group has non-circular preference between whole states. This fact shows that, with proportional voting, the CDMP of Equation 6 will terminate with a final proposal (the CP) rather than just wandering. Similarly, the CPT suggests that with proportional voting, the CDMP of Equation 9 will stabilize and terminate, though usually without all actors occupying the same position.

Our tax/subsidy example simulates the CDMP of negotiating through the space of possible states to reach the central one. As mentioned earlier, this could potentially be a very difficult optimization, depending on the problem structure. In CGE models for carbon ETS policy, one might take the set of policy positions to be discrete. With 20 regions that might or might not participate in the ETS, there are 2^{20} or about a million ways to choose which are or are not included, and each choice might have very different impacts on different regions. There could be further discrete choices of regulations governing the market structure, e.g. whether there was one national market, separate markets for each regions, separate markets for each industrial sector, etc. In our example integration, we have designed the set Ω of policy options so that simple hill-climbing performs well for this example.

Strategic voting allows actors to assess not only the intrinsic desirability to themselves of a position but also the support or opposition it is likely to get from other actors. In an economic context, it would always be desirable to an individual actor that all the other actors give enormous subsidies – but such a position would not win any support. Advocating a shortsightedly self-interested policy would likely fail and hence would have lower expected value than advocating a policy of enlightened self-interest which compromised just enough to win strong support from other actors with aligned interests.

Strategic voting can give counter-intuitive behavior, such as supporting positon X even though it was less desirable than the status quo Y, because doing so split the coalition supporting Y and increases the probability that some other option Z might win, where Z was more desirable than Y. In other words, supporting X increases the expected value of the whole state, even though the position X was less desirable than the position Y.

3 Simple Leontief Economy

So far, we have explained how a PCE generalizes the standard academic framework for negotiation so that we can address problems with more complex DSUM, e.g. economic models. We now describe a very simple model of an export-driven Leontief economy; it is implemented as the demonstration program `leonApp` in the KTAB distribution. The model first generates synthetic economic data, derives the parameters of the Leontief model from it, then runs two different CDMP on the same data set. The source code, and the data for the particular example discussed here, are on the KTAB website.

More interesting policy questions would require more complex economic models, but even this simple model has several features which are difficult to handle with more traditional policy models. When analyzing the question of what political-economic factors are likely to influence the scale and structure of a potential carbon ETS, the same PCE-based framework could be integrated with (e.g.) a multi-regional CGE.

In this example, each actor represents an economic interest group. There are N=5 industrial sectors, each with a corresponding actor that represents trade associations, lobbyists, and similar groups advancing the interest of that industry. The model does not specify how or through what channel these interests are expressed and influence exerted. The actors could be industrial leaders personally lobbying a powerful central executive; they could be dispersed collections of individuals in public meetings with their local government representatives. The model is concerned with the existence of a group and the level of influence exerted, not the means or institutional structures through which it is exerted.

There are M=2 value-added factors of production, notionally unskilled labor, skilled labor, and capital. Each has an associated actor, representing the common economic interests of those groups, as expressed through labor unions, elected representatives, etc. Again, many political analyses focus on “framing the issue” and similar techniques for swaying interest groups, but this model is concerned with the level of influence exerted, not the techniques of influence.

Finally, the synthetic economic data also includes K=2 different consumption groups, notionally household and government. We will not present more detail because the consumption groups play no significant role in this CDMP.

There are natural economic linkages between different industrial sectors, which can result in their promoting a common policy on taxes and subsidies. Because the value-added is counted for both sectors and factors, there are natural economic linkages between them: an increase in value added for unskilled labor used in the agricultural sector would be a benefit to the agricultural sector in general and to the unskilled labor factor in general.

The position of each actor is a vector τ of taxes and subsidies on N export goods. The base-year export quantities are the vector X_0 and the base-year prices are P_0. For the i-th export commodity, the base-year price is $p_{i0} = 1$, as is standard in CGE models. The vector elements have $\tau_i > 0$ for a tax and $\tau_i < 0$ for a subsidy. After applying the tax/subsidy vector, the price of the i-th export commodity is the following:

$$p_{i\tau} = 1 + \tau_i$$

The vector of after-tax prices is P_τ with corresponding after-tax export demands X_τ. Export demand is constant-elasticity in price for each commodity separately:

$$x_{i\tau} = x_{i0}\left(\frac{p_{i0}}{p_{i\tau}}\right)^{\varepsilon_i}$$

The tax/subsidy vector is constrained so that after-tax price is no less than 1/10-th and no more than 10 times the base-year price. It is further constrained to be revenue-neutral and self-financing after demand is adjusted for the new prices. This is a simple way to avoid the unrealistic solution of giving every actor the maximum possible subsidy. The set Ω of possible tax positions is therefore the following:

$$\Omega = \{\tau \mid 0 = \tau \cdot X_\tau, -0.9 \leq \tau \leq 9\}$$

Given a linear sub-model of how wages, interest, and government revenues are distributed to the K consumption groups, and the Cobb-Douglas utility functions of the consumption groups, we can use basic linear algebra to calculate the N-by-N matrix, A, which specifies how total production depends on exports:

$$Q = AQ + X$$

$$Q = (I - A)^{-1}X = L_\alpha X$$

With M value-added factors of production and N industrial sectors, the basic Leontief model says that the amount of value added for each factor, f, per unit of production in a sector, s, is a constant, ρ_{fs}. Thus, M-by-N matrix of added value, V, is a linear function of the N-by-1 vector of production. This can be represented by a simple multiplication:

$$V_{fs} = \rho_{fs}Q_s$$

We will abbreviate this component-wise multiplication to $V = \rho Q$, where the indices are assumed.

As mentioned, the factors and sectors use different L matrices to estimate Q from the same X. This is the economic model used by the factors to estimate the effect on them of export tax policies:

$$V^F(\tau) = \rho(L_\alpha X_\tau)$$

Note that this V^F is an M-by-N matrix. The value added expected by factor f is the sum over sectors for that factor:

$$v_f^F(\tau) = \sum_{s=1}^{N} V_{fs}^F(\tau)$$

As stated earlier, it is possible to extend the model to take into account the investment needed to support a constant, balanced rate of growth$g$. Again, our goal is to demonstrate integration of bargaining and economic models, rather than to reproduce sophisticated economic growth models such as the Solow growth model or DSGE. We treat growth as simply a constant percentage increase in output and capital stock. This changes to fundamental equation to include demand for investment goods:

$$Q = X + AQ + N$$

In any given year, the capital required to produce all the output cannot exceed the available capital stock:

$$BQ \leq K$$

If we assume that the rate of depreciation is constant across industries, we get the required vector of investment in each period required to support steady growth:

$$K_{t+1} = (1 - d)K_t + N_t = (1 + g)K_t$$

This shows that the required investment is the proportion required to cancel depreciation and provide the desired growth:

$$N_t = (d + g)K_t$$

Assuming no excess capital stock means that$BQ_t = K_t$, so we can get the required level of investment:

$$N_t = (d + g)BQ_t$$

This leads to the revised estimate of output, given export demand and expected investment:

$$Q = X + AQ + (d + g)BQ$$

$$Q = \big(I - (A + (d + g)B)\big)^{-1} X = L_\beta X$$

This is the economic model used by the sectors to estimate the effect on them of export tax policies, using the same ρ matrix as before:

$$V^S(\tau) = \rho\big(L_\beta X_\tau\big)$$

Again, this V^S is an M-by-N matrix of value added. The value added expected by sector s is the sum over factors for that sector:

$$v^S{}_s(\tau) = \sum_{f=1}^{M} V_{fs}^S(\tau)$$

The combination of these two simple economic models form the DSUM for this example.

It is significant that the CDMP does not assume that all the actors agree on the correct economic model. This is allowable because we are not trying to model what the actual consequences of a policy will be; we are trying to model how the CDMP will proceed, even when actors disagree about possible future consequences. Where they agree that a policy increases overall GDP, without loss to any actor, they would probably all favor implementing such a Pareto-improvement[1]. However, the criterion of Pareto-optimality merely directs actors to the efficient frontier; it does not determine which particular Pareto optimum will be chosen. In other words, the actors will increase GDP where they can (and where they agree that the policy has that effect), but they will bargain over how to distribute the gains. Because they do not necessarily use exactly the same economic model, the actors might not agree on the precise location of the GDP-maximizing Pareto improvements. Indeed, it is perfectly possible that some actors might use a very complex DSUM to estimate the consequences of a policy, while others might use a very simple one. One classic example is the contrast between those actors whose only economic concern is the level of a legally mandated minimum wage and those actors who look at complex pricing and substitution effects. The fact that different actors can have very different ideas of the consequences of a policy helps explain how the CDMP can settle on policies which are not Pareto optimal according to any of the actors involved.

To assess strategically viable policy options in a particular real economy, much more complete and realistic DSUM, combining economic and policy models, would obviously be required. A minimal economic model might be CGE with validated base-year data. To address disparate regional impacts, an IRIO or MRIO might be appropriate. To address negotiation over rationing, rents, quotas, and so on, a mixed-complementarity model might be appropriate. A minimal policy model might be to identify the interests of actors directly with particular economic actors, as is done in the `leonApp`. A more sophisticated model might take the weighted combination of several sectors, as appropriate to the organizational networks or geographic districts supporting particular factions. Another variation would be to model an actor with concerns about inequality by making their utility of a policy be increasing with mean value added but decreasing with variance: more wealth is desirable, but high inequality is undesirable.

One model of bargaining over taxes and subsidies would be to choose them so as to maximize GDP, then use the benevolent hand to redistribute the benefits. While standard, this approach glosses over some fundamental problems which our approach is designed to address (albeit in a simplified manner).

First, there are many Pareto-optimal ways to redistribute wealth, including no redistribution at all. The "maximize then redistribute" approach simply shifts the bargaining problem to the question of how to redistribute. Obviously, there are some groups who gain by redistribution and some who lose, so there will be bargaining over how to do so. When long-term effects are taken into account, it is entirely possible that different actors will have different discount rates and time preferences, so the issue becomes one of inter-temporal redistribution – which again is a problem of bargaining among Pareto optimal outcomes.

---

[1] Some actors might oppose it on the grounds that it increased inequality. The converse situation, where a desire to increase equality led to policies which gave a general Pareto worsening, is not historically unprecedented.

Second, there is disagreement about which policies will maximize GDP. Again, one classic example is the debate over the effect of minimum wage laws, where powerful interest groups disagree as to the net effect of such policies.

Third, there is disagreement about the redistribution effects of different policies. The BFM directly addresses redistribution via bargaining, but it includes the assumption that policy makers can precisely target the redistribution with no constraints due to economic linkages. As soon as those linkages are taken into account, not only do we have a more complex economic model than BFM can handle but also there is the possibility that actors will disagree as to the effects of redistribution policies.

Fourth, as is quite well-known, the decisions to redistribute wealth change the incentives to produce and to consume. Hence, the first step (maximize GDP) cannot be analyzed separately from the second step (redistribute wealth), lest the redistribution undermine the production and result in a much lower GDP than expected. Like other approaches, this paper's approach considers both in an integrated fashion.

An advantage of this approach is that the formal framework of negotiation is not tied to a fine-tuned mathematical model of policy options or utility functions, so it can remain unchanged as the DSUM becomes more complex and realistic; research is ongoing to integrate these more sophisticated DSUM into several CDMPs using the KTAB toolkit.

However, this simple combination of bargaining with a Leontief economy, as implemented, already has some essential features which are not represented in SMP or BFM:

- The number of policy parameters is neither fixed at 1 (vs. SMP-1) nor fixed at the number of actors (vs. BFM). It can be varied as necessary to represent the "policy levers" actually available to actors in the situation of interest.
- The value of a policy to an actor depends not on the Euclidean distance between policies but on the variation in economic outcomes of those policies (vs. SMP). Because of threshold effects and non-monotonic responses, the variation between outcomes might be very different from the distance between policies.
- The presence of economic interactions can create natural alliances between actors (unlike both SMP and BFM), so that they tend to benefit or lose together, regardless of their strength.
- Because of economic interactions, actors cannot design policies to precisely assign whatever benefits they choose (vs. BFM). Technically, policies do not have enough degrees of freedom to separately control each actor's utility, so that an actor may not be able to craft a policy which helps precisely those they want to help. They may be forced to choose which of several allies to cultivate, rather than craft a policy which wins support from all. Similarly, they may be forced to help a weak actor whom they do not strategically need, because they do need a strong ally whose utilities happen to be correlated with the weak one. This never happens in the BFM: if a weak actor is not part of any minimum winning coalition, then no other actor has any incentive to help them and they never receive anything in the BFM.

4 UTILITIES

The utility to each actor from a position is the value-added which they expect from their particular economic model under that policy, normalized into a [0, 1] scale. The normalization is based on Monte Carlo sampling. The first step is to generate 2500 random (but revenue-neutral) policies and record the value added expected by each sector or factor, using their respective economic model. Many psychological studies have shown that people tend to view potential losses of what they have as more significant than potential gains. This effect is modeled by converting the value-added numbers to "raw utility" values: losses compared to the base year are weighted heavier than gains, where we chose $r = \frac{3}{2}$ for this example:

$$u_i'(\tau) = \begin{cases} \nu_i(\tau) \text{ if } \nu_i(\tau) \geq \nu_i(0) \\ r\big(\nu_i(\tau) - \nu_i(0)\big) + \nu_i(0) \text{ if } \nu_i(\tau) < \nu_i(0) \end{cases}$$

Separately, each actor linearly rescales the raw utilities so that the position with their worst plausible outcome gets utility 0 and their best plausible outcome gets utility 1. For each actor, the utility function as a function of value-added is concave and piecewise linear.

While the utility to each actor is piecewise linear as a function of value-added, the value-added itself is not a linear function of the tax-policy vector. As argued earlier, it is not even single-peaked. However, it is still possible to verify that in this example, the value-added for each sector does show the expected decline as the export tax on that sector increases. To demonstrate the point, we present a table of the correlation of the value added to a sector, $\nu_i(\tau)$, versus the export tax on that sector, τ_i, over the 2500 samples. Thus, we are looking at the correlation of that actor's value added with just the corresponding component of the tax policy.

Sector	Correlation
2	-0.9624
3	-0.9629
4	-0.4182
5	-0.9362
6	-0.8805

Table 2: Correlation of sectoral tax rates with sectoral value added

It is not a perfect linear correlation for several reasons:

- The export demand response to a tax rate is nonlinear, each to differing degrees.
- Total value added in a sector depends not only on its own direct export demand, but also on the domestic demand indirectly generated by export from other sectors. Those export demands depend on the other components of the tax vector
- Because the policy is revenue neutral, taxing one sector tends to provide subsidies to the other sectors (which have their own nonlinear demand curves): taxing one export sector reduces export demand but may drive up domestic demand. For this data set, the effect is particularly strong for sector 4.
- Monte Carlo sampling error, though that is small when 2500 samples are used.

5 EXAMPLE CDMPS FOR EXPORT TAXES AND SUBSIDIES

With the general KTAB framework in place, the integrated policy-economic model is quite simple to explain. The model is implemented as the command-line demonstration program `leonApp` in the KTAB toolkit; which CDMP is used is controlled by a command-line input. They can both be used on the same model, so as to contrast their results when applied to the same economic model.

To initialize the model with synthetic data, `leonApp` uses a random-number generator to synthesize plausible data for the Leontief model: A, B, ϵ, X_0, ρ, g, and d. This is done by generating a full table of base-year economic transactions, then deriving the parameters in the same manner as do most CGE models. The number of factors, the number of sectors, and the number of consumption groups are variable parameters.

The capability of an actor to exert influence is taken to be their value-added in the base year, somewhat analogous to Negeshi weights [Negeshi 1972]. While there are well-known normative problems with Negeshi weights, the particular method of determining influence is not essential to the demonstration of integration between policy-making and economic-models.

5.1 OFFICE-SEEKING POLITICIANS

In the first CDMP, we model the behavior of office-seeking politicians (OSP) proposing policies. There is one current proposal, z_t, and a sequence of amendments are offered as per Equation 6. We use proportional voting so that the process will terminate at the Central Position.

Notice that in this process, the OSP actors are separate from the N sectors and the M factors. Just as with Black 1948, there is actually no need to explicitly model the characteristics of the OSP, because their actions are designed solely to maximize the net support from the N+M actors who are explicitly modeled.

In terms of implementation, we simply perform one hill-climbing search starting with a status-quo proposal z_0 of zero taxes on all goods. The search through potential positions is done by a simple hill-climbing search algorithm. At each step of the search, the neighborhood considered is the set of positions reachable by a line search that varies one or two rates. The hill-climbing search takes the best new vector it finds at each step of the way, until it reaches a local maximum, which in this case is the CP. Notice that each step of search can be considered a simulation of the OSP actors as they do unilateral changes or bilateral deals so as to win allies and build an advantageous coalition; one hill-climbing search simulates the whole CDMP. Just as in the MVT, the competition between OSP actors drives the group to the Condorcet Winner, which in this case is the CP.

For this model, we assume that all the OSPs are responsive to the same set of constituents, i.e. the factors and sectors.

For this data set, the CP is the following vector. Because the first two actors, #0 and #1, are the factors of production, the sectors actually start at #2. Sector #4 is most subsidized, sector #6 is less subsidized, and the other three sectors are taxed.

Tax on Sector 2	Tax on Sector 3	Tax on Sector 4	Tax on Sector 5	Tax on Sector 6
+0.0025	+0.0043	-0.2290	+0.0658	-0.1187

Table 3: Central Position reached by Office-Seeking Politicians

There are two interesting uses of this CP. First, it is of course an estimate of the most likely consensus position which OSP would adopt. Second, it is useful to see how each actor's utility-maximizing position differ from the OSP position. As mentioned earlier, the OSP position is well-approximated as a compromise between the actors' self-interested positions, because the actors all have concave (compromising) utility functions. If they were convex (uncompromising) then the OSP would tend to jump among the extremes without compromising between them.

5.2 Utility-Maximizing Actors

In the second CDMP, we model the behavior of actors proposing self-interested policies. Each actor has their own separate position at each turn of the CDMP, and they modify them in parallel as per Equation 9.

As mentioned, the value to an actor of a state is the expected value of that state, using a PCE with proportional naïve voting by actors within a potential state. While the actor could advocate a position highly advantageous to itself, the expected value of the state may be quite low because no other actor supports the position. Therefore, optimizing one's position to maximize the expected value of the resulting state will automatically seek out the best balance (in the current situation) of pushing one's self-interest and offering inducements to potential allies. Clearly, a very powerful actor would not need to make as many concessions to allies as would a weak actor.

At the start of the CDMP, every actor has the same status-quo initial position: zero taxes on all goods.

As with the MVT, CPT, and BFM, the CDMP proceeds over several turns, until all the positions stabilize. Because of the CPT, the preferences are non-circular so the process will stabilize, though the actors will generally not all stabilize at the same position. At each turn, each actor separately finds a position for themselves which maximizes the expected value to them of the resulting state. In other words, they look at the expected value of hypothetical states in which they change their own position while assuming all other actors do not change their positions. The search through potential positions is done by a simple hill-climbing search algorithm. The hill-climbing search takes the best new vector it finds at each step of the way, until it reaches a local maximum, which is the actor's new position for the next turn. Thus, each actor separately performs a hill-climbing search, for each step of the CDMP. In terms of Equation 9, the neighborhood explored at each turn is the entire spaceΩ. Notice that each step of search can be considered a simulation of that particular actor's backroom negotiation so as to win allies and build an advantageous coalition.

In each turn of the CDMP, every actor independently locates their optimum position assuming others do not change. But because other actors do adjust their positions, the turns continue until no actor changes significantly, i.e. until a Nash Equilibrium is reached.

With 2 factors, 5 sectors, and 2 consumption groups, the CDMP usually completes in two or three turns, sometimes five to seven, depending on what synthetic data is randomly generated. With 5 factors, 10 sectors, and 5 consumption groups, the CDMP usually completes in 4 or 5 turns.

We do not expect to find a succinct analytical description of the Nash Equilibria unless more structure is specified in the economic model.

Sector	Tax on Sector 2	Tax on Sector 3	Tax on Sector 4	Tax on Sector 5	Tax on Sector 6
2	**-0.0795**	+0.0664	-0.0453	+0.1121	-0.0686
3	+0.0343	**-0.0473**	-0.0476	+0.0940	-0.0873
4	+0.0020	+0.0021	**-0.3093**	+0.0758	-0.1051
5	+0.0143	+0.0045	-0.0700	**+0.0110**	-0.0938
6	+0.0049	+0.0081	-0.1084	+0.1083	**-0.1821**

Table 4: Positions reached by Utility-Maximizing Actors

5.3 Contrasting Office Seeking and Utility Maximizing Outcomes

The results differ – as expected even in one dimensional case – because they represent fundamentally different CDMPs. Comparing the emboldened diagonal values to the OSP position, we can see that each sector's equilibrium position is more favorable to itself than is the OSP position, as expected.

If we average down each column, we get the average of the self-interested positions. As expected, it is quite similar to the OSP position. Treating them as vectors, the correlation of the average position and the OSP position is +0.9246.

Tax on Sector 2	Tax on Sector 3	Tax on Sector 4	Tax on Sector 5	Tax on Sector 6
-0.0006	+0.0046	-0.1048	+ 0.0721	-0.1109

Table 5: Mean Position reached by Utility-Maximizing Actors

The general pattern observed for test data is that the actors adopt similar overall positions, as they are all driven to make concessions to the most powerful interests. This is shown in the following table, which shows the correlation of each actor's preferred position with the OSP position.

Sector	Correlation
2	+0.6151
3	+0.7029
4	+0.9864
5	+0.8521
6	+0.8308

Table 6: Correlation of Actor Positions with OSP Position

However, when we examine the correlations in how individual actors differ from the mean position, $\bar{\theta}$, we see a more complex pattern. Those actors with positive economic linkages tend to gain (or lose) together by subsidizing (or taxing) the same set of goods, so they would like to make similar changes from the mean position. The degree of common (or opposing) interests can be indicated by the following matrix of correlations:

$$\rho_{ij} = \rho(\theta_i - \bar{\theta}, \theta_j - \bar{\theta})$$

Note that actors 0 and 1 are the two factors of production, while the other five actors represent sectors. A positive correlation means that they would both like to change from the average position in roughly the same way; a negative correlation means that they would like to change it in opposite ways. To avoid confusing symmetries, the diagonal and lower triangle are greyed-out. The absolute values greater than 0.5 are emboldened.

ρ_{ij}	0	1	2	3	4	5	6
0	+1.0000	-0.1082	-0.3005	-0.4169	-0.1158	+0.3912	+0.3884
1	-0.1082	+1.0000	+0.0668	**+0.8669**	**-0.8100**	**+0.6527**	-0.3567
2	-0.3005	+0.0668	+1.0000	-0.0585	-0.4556	-0.0862	-0.1914
3	-0.4169	+0.8669	-0.0585	+1.0000	**-0.6068**	+0.2355	-0.1486
4	-0.1158	-0.8100	-0.4556	-0.6068	+1.0000	-0.4764	+0.0282
5	+0.3912	+0.6527	-0.0862	+0.2355	-0.4764	+1.0000	**-0.5846**
6	+0.3884	-0.3567	-0.1914	-0.1486	+0.0282	-0.5846	+1.0000

Table 7: Correlations of Changes Desired by Actors

The strongest common interests are between factor 1 and sector 3 (+0.87) and between factor 1 and sector 5 (+0.65). Because of the revenue-neutrality constraint, the sectors tend to have clashing interests (the subsidy one receives must be provided by another), but still sector 3 and sector 5 have some overall common interest (+0.24). The sharpest conflicts in the table are between sectors 4 and 1 (-0.81), between sectors 4 and 3 (-0.61), and between sectors 5 and 6 (-0.58).

While this example uses only synthetic data, it does illustrate some of the results which can be obtained and conclusions which can be drawn.

6 Conclusions

We have reviewed the basic concepts of bargaining models and how they differ from traditional economic models. The strengths and limitations of the median voter theorem and the Baron-Ferejon models have been discussed, particularly the problem of being unable to accommodate plausible economic models with a realistically limited number of control parameters. The central position theorem and probabilistic Condorcet elections offer two distinct ways to extend the bargaining paradigm. Each models a different kind of collective decision making process. Not only are the two CDMPs applied to the same data set, but the contrast between the two kinds of results provides further insight into the likely pattern of alliance (and opposition) between different economic actors for this tax policy.

The next phase of research will focus on three main areas:

- Applications
- Algorithmic Efficiency

- Analysis

The most important research area is to identify a current policy problem which would benefit from this approach. Some candidates being considered include the selection of regions to participate in a carbon emission trading scheme (ETS), design of the ETS, subsidization of energy products in domestic markets, and policies for achieving local content goals.

With realistic problems come large data sets, high-dimensional parameterization, and a combinatorial explosion of policy designs. Meeting these challenges will require much more efficient solution algorithms. While there are many plausible approaches to increasing efficiency, we will focus on two. The first and more general approach is to implement cloud-based parallel computing. In the CDMP for actors maximizing their expected value, each actor's search process could be launched in a separate, parallel computation. In both CDMPs, the policies in a neighborhood "near" the current best solution could also be evaluated in parallel. The second and less general approach is to design specific solution algorithms tuned to the structure of the problem being considered, analogous to replacing the generic simplex algorithm by the transshipment algorithm. How the generic approach should be modified would naturally depend on the particular problem structure being addressed.

The analytical characterization of the CDMP has two sub-tasks. The bargaining process of the office-seeking politicians is definitely known to converge because of the CPT. However, the CDMP that actors use to maximize expected value of a state has not been thoroughly analyzed for convergence and stability. Nor have the Nash Equilibrium positions been adequately analyzed and characterized.

7 References

Arrow, K.J., "A Difficulty in the Concept of Social Welfare", Journal of Political Economy 58(4) (August, 1950), pp. 328–346

Black, Duncan. 1948. "On the Rationale of Group Decision-making." The Journal of Political Economy 56 (1):23 - 34.

Condorcet, Marie Jean Antoine Nicolas de Caritat, marquis de, 1785, “Essay on the Application of Analysis to the Probability of Majority Decisions”

Coughlin, P. (1992) *Probabilistic Voting Theory*, Cambridge University Press.

Feder, S. (2002). “Forecasting for Policy Making in the Post Cold War Period”. Annual Review of Political Science 5: 111-125

Negishi, T. (1972) *General Equilibrium Theory and International Trade*, North-Holland Publishing Company

Poole K. and Rosenthal H. (1985) “A Spatial Model for Legislative Roll Call Analysis”. American Journal of Political Science 35: 228-278

Poole K. and Rosenthal H. (1991) “Patterns of Congressional Voting”. American Journal of Political Science 29: 357-384

Poole K. and Rosenthal H. (1999) “D-Nominate after 10 Years: A Comparative Update To Congress: A Political Economic History of Roll Call Voting”. Available at SSRN: http://ssrn.com/abstract=1154162

Poole K. and Rosenthal H. (2000) *Congress: A Political-economic History of Roll Call Voting*. Oxford University Press

Wise, (2010a). “Technical Overview of Lodestone”. RAND (Unpublished Draft).

Wise, B (2010b). “Utility, Voting and Strategy in General Issue Spaces”. RAND (Unpublished Draft).