Pairwise Majority Rule with Multi–dimensional Choices

1 A Geometric Example

Although single–peaked preferences, or the single–crossing property, guarantee that there will be no cycling, these properties are defined only when the set of alternatives is "one–dimensional". *Persson and Tabellini* mention a property which will guarantee the existence of a Condorcet winner when the set of alternatives is multi–dimensional : the "intermediate preferences" property defined in Definition 4 (pg. 25). But this property is a pretty special one — it requires that people can somehow be described by a single–dimensional characteristic. And even then, intermediate preferences is a pretty special requirement.

Here I want to illustrate a little buit the problems which may arise when the set of policy alternatives is multi–dimensional, even when the preferences of voters are quite well–behaved. What I am trying to show here is that the natural extension of single–peakedness to more than one dimension won't be good enough.

To see this, suppose that there were two variables which voters had to choose. For example, suppose X is spending on police services, and Y is the level of spending on education. The cost of these services will be shared by all the voters. (So a point (X, Y), in which both X and Y were very high, would represent a policy of spending a lot on both police services and on education, and levying high taxes to pay for all this expenditure.)

Suppose that each voter ranks alternatives by how close the alternative is to her preferred policy. For example, suppose that voter #1 had a preferred policy of $(700, 200)$: this voter wants to spend a lot on police services but not much on education. So $(700, 200)$ is her most preferred policy. Other policies she ranks by how far they are, in distance, from her preferred policy. A policy $(400, 500)$ is closer to $(700, 200)$ than is $(300, 600)$, so she prefers $(400, 500)$ to $(300, 600)$. Her preferences could be represented by a utility function

$$U^1(X, Y) = 300000 - (X - 700)^2 - (Y - 200)^2$$

since the distance of any policy (X, Y) from her most preferred policy $(700, 200)$ can be measured as

$$d = \sqrt{(X - 700)^2 + (Y - 200)^2}$$

So she is indifferent among all policies which are the same distance from $(700, 200)$. In two dimensions, her indifference curves are circles around the point $(700, 200)$, as illustrated in figure 5. Policies inside an "indifference circle" are preferred to policies outside the circle.

2 Ranking Different Policies

The overall utility from different policies can be graphed in 3 dimensions, as in figure 6. This figure shows that her preferences look sort of "single–peaked" : her utility graph has a peak at $(700, 200)$ (at which $U^1 = 300000$), and it falls off in every direction. In the graph, $(700, 200)$ is the only peak.

Her preferences also look single–peaked in any single dimension. Figure 7 shows her utility, as a function of X, holding Y constant (at $Y = 500$). Her utility reaches a peak at $X = 700$, and falls off in either direction, as we move left or right from $X = 700$.

Imagine now that there are several people, with similar sorts of preferences, differing only in their most preferred policy. For example, suppose that there are 3 voters, with

$$U^1(X, Y) = 300000 - (X - 700)^2 - (Y - 200)^2$$

$$U^2(X, Y) = 300000 - (X - 300)^2 - (Y - 300)^2$$

$$U^3(X, Y) = 300000 - (X - 200)^2 + (Y - 800)^2$$

so that person 1's preferred policy is $(700, 200)$, person 2's is $(300, 300)$ and person 3's is $(200, 800)$. Each person's 3-dimensional utility graph looks like figure 6, except that the peak is at a different point for each person. Person 2 wants relatively low spending on both categories of expenditure, and person 3 wants a lot of educaton spending, and very little spending on police services.

So these preferences seem "just like" single–peaked preferences, in that moving further away from a preferred policy (in any direction) moves the person to lower and lower levels of utility.

3 One Dimension at a Time

And if we voted (by pairwise vote) on these issues one dimension at a time, there would be a winner in every dimension. For instance, consider voting about the level of police expenditure X. Each person's utility from a given level of X is

$$U^i(X, Y) = 300000 - (X - X_i^*)^2 - (Y - Y_i^*)^2$$

where (X_i^*, Y_i^*) is the person's preferred policy (for example $(300, 300)$ for person #2). Differentiating,

$$\frac{\partial U^i}{\partial X} = -2(X - X_i^*)$$

so that utility increases with X whenever $X < X_i^*$, reaches a peak at $X = X_i^*$, and then falls as X increases above X_i^*, just as shown in figure 7.

So, with these three people, if they all voted sincerely, and if we voted one issue at a time, then there would be a Condorcet winner in each dimension : $X = 300$ and $Y = 300$, in each case the median of the three people's preferred levels of X and Y.

4 Many Dimensions at Once

But restricting the vote to one dimension at a time is arbitrarily controlling the agenda. What if a person could propose changing — simultaneously — both X and Y? That is, if $(X, Y) = (300, 300)$ is the status quo, what if someone could introduce a new bill, proposing a totally new (X, Y) combination, in which $X \neq 300$ **and** in which $Y \neq 300$?

It turns out that $(300, 300)$ can be defeated — even though $X = 300$ and $Y = 300$ are winners if we can only make changes in a single dimension at a time.

What if someone proposed $(400, 400)$ as an alternative to $(300, 300)$? Person 2 would obviously vote against the proposal, since $(300, 300)$ is his most–preferred policy. But person 1's utility from $(400, 400)$ is 170,000, which is higher than the utility of 130,000 which she gets from $(300, 300)$. And person 3's utility from $(400, 400)$ is 100,000, which is higher than the utility of 40,000 which she gets from $(300, 300)$. In other words $(400, 400)$ is closer than $(300, 300)$ to both person 1's preferred policy $(700, 200)$, and to person 3's preferred policy $(200, 800)$.

So $(300, 300)$ is not a Condorcet winner, since it gets defeated by $(400, 400)$. The policy $(400, 400)$ is also not a winner. If it were the status quo, someone could propose reducing X to 300 — holding Y constant at 400 —, and people 2 and 3 would vote for such a reduction.

In fact, there can be no winner in this example. It was just shown that $(300, 300)$ could be defeated by $(400, 400)$. And any other policy, in which $X \neq 300$ or in which $Y \neq 300$, could itself be defeated, either by changing X to 300, or by changing Y to 300.

The problem? Even though preferences look nice and convex, and even though the graph of utility looks single–peaked if we slice it in any direction, there is no way in which the alternatives can be lined up in a single dimension. If choices are inherently multi–dimensional, and if there are no restrictions on the agenda, then there will be no overall winning policy under pairwise majority rule.

So one solution seems to be to vote on policies one issue at a time. The committee system in legislatures seems to achieve this sort of effect : in the committee on police services, only changes in X can be proposed, not changes in Y.

5 How to Slice Policies

But there may not be a natural way to divide the issue into "dimensions". One variable that has been changing in the background in all these examples is the tax rate. Presumably increasing spending on police services, or on education, will serve to increase taxes. That is why each of these voters does not want an infinitely high level of spending on any category : taxes go up when spending goes up.

As described so far, voters choose expenditure on different categories, with the tax rate adjusting "invisibly" so as to pay for the expenditure. But that is not the only way that legislatures work. Often, the total amount of taxes is decided (perhaps by some committee). So voters may choose a level of total expenditure $E = X + Y$ in one committee. In this committee the choice is one–dimensional : the level of expenditure (which equals the total taxes levied). In some other committee, they then choose how the given level of expenditure E is to be divided between police services and expenditures.

In other words, this last paragraph has proposed a new set of restrictions on how policies can be changed. First, looking at expenditure one category at a time, new proposals were restricted to changes in X, or changes in Y, but not both. If someone proposed changing X, Y was not changed. Now someone (in the tax setting committee) can propose a change in total expenditure E, but **not** in how the expenditure is divided between education and police services. In the other committee, someone can propose changing how expenditure is split between police services and education, but **not** the total level of expenditure. That is, new proposals are now restricted to one of : "increase or decrease expenditure on each category (police and education) by some percent (but the percent has to be the same)", or "increase (or decrease) the share of the given total expenditure which goes to education (but don't change the total expenditure)".

Any restriction on how much can be changed by a single proposal will have the same effect : it makes each choice effectively one–dimensional, and gets rid of the cycling. But how the restrictions get imposed does affect the policy which will actually get chosen. When cycling is possible (without any restrictions on new legislation), the details of the rules of the legislature will wind up determining the policy that wins.