## Public Choice : ($c$) Single–Peaked Preferences and the Median Voter Theorem

The problem with pairwise majority rule as a choice mechanism, is that it does not always produce a winner. What is meant by a"a winner"? If voters are choosing among a group of alternatives $A, B, C, \ldots$, using pairwise majority rule, a policy $W$ is said to be a **Condorcet winner**, if $W$ defeats any other policy in a pairwise vote. That is, $W$ is a Condorcet winner only if at least half the voters prefer $W$ to $A$, at least half the voters prefer $W$ to $B$, and so on.

The "standard" example of cycling, in which there are three people, and three alternatives, $A, B, C$, and preferences are

| 1 | 2 | 3 |
|---|---|---|
| $A$ | $B$ | $C$ |
| $B$ | $C$ | $A$ |
| $C$ | $A$ | $B$ |

shows that, in some situations, there maybe no Condorcet winner. But this problem of cycling does not arise in all situations.

So one way of dealing with the problems shown by Arrow's Impossibility Theorem is to look for situations in which the cycling problem does not arise under pairwise majority rule. Pairwise majority rule obeys three of Arrow's axions : $P$, $ND$, and $IIA$. Maybe it is a good idea to give up on $UD$ : to consider only environments in which people's preferences are lined up in such a way as to ensure that cycling will not be a problem.

It turns out that cycling will not be a problem, if people's preferences are aligned in a particular way : if preferences are "single–peaked".

Preferences are said to be **single–peaked** if there is some way of lining up all the alternatives, so that the graph of **every** voter's preferences has a single local maximum. The graph of a voter's preferences lists all the alternatives along the horizontal axis, and then on the vertical axis has the "score" assigned by the voter to each alternative.

For example, in figure 1, the alternatives are graphed along the horizontal axis. On the vertical axis, some measure of the voters' utility is graphed : higher scores mean an alternative is preferred by the voters. Voter 1's most–preferred policy in the figure is $D$, follwed by $C$, then $E$, then $B$ and finally $A$. Voter 4 ranks $A$ first, followed by $B$, $C$, $D$ and $E$ in that order.

A voter's preferences are said to be single–peaked if the graph (with the alternatives along the horizontal, and the scores on the vertical) has at most one peak : either it goes up and then down, or it keeps going up as we move left to right, or it keeps going down.

In figure 1, all five of the voters have single–peaked preferences : that is, either the scores go up until a maximum is reached, and then down, or the scores go straight up ( like voter 2's ) or straight down (like voter 4's ). But in figure 2, voters 3 and 5 do **not** have single–peaked preferences. Voter 3, for example, has a local maximum at alternative $B$ and another local maximum at alternative $D$.

So preferences are single–peaked whenever there is **some** way of arranging the alternatives along a line so that the graphs of each voter's preferences have just one local maximum.

Under pairwise majority rule, if voters' preferences are single–peaked — that is, if there is some way of arranging the alternatives in a line, left to right, so that the graphs of the preferences look like those in figure 1 — then majority rule will generate a transitive ordering, and a Condorcet winner.

In particular :

**Median Voter Theoem** : If there is some way of lining up the alternatives on a graph so that all voters' preferences are single–peaked, then the **median** of the voters' most–preferred alternatives will defeat any other alternative in a pairwise vote.

The median of people's preferred alternative is the alternative which has the number of people who prefer an alternative to the left equal to the number of people who prefer an alternative to the right.

[The concept is the same as, for example, the median income among a group of people. The median income is a level of income such that there are equal numbers of people with higher income, and with lower income. If five people had incomes (in thousands of dollars per year) of 10, 12, 15, 19 and 44, then the median income would be 15 : two people have higher income, and two people have lower income. The median income is different from the **mean**, or average income, which is not 15 in this case, but 20.]

In the situation depicted in figure 1, it is alternative $B$ which is the winner. It is the median of the voters' most–preferred alternatives: 2 voters (#'s 1 and 2 ) have a preferred choice to the right of $B$, and 1 voter ( voter #4) has a preferred choice to the left. So if alternative $C$, $D$ or $E$ is proposed, it will lose in a vote to alternative $B$ : voters #2, 3 and 5 all prefer alternative $B$ to anything to the right of $B$. If alternative $A$ is proposed, then it loses to alternative $B$ : everyone except voter #4 prefers alternative $B$ to alternative $A$.

So, at least in the example of figure 1, there is a Condorcet winner. Alternative $B$ will defeat each of the other 4 alternatives in a pairwise vote. Because $B$ is the median of the most–preferred alternatives, a majority of the voters prefer $B$ to anything to the right of $B$, and a majority of voters prefer $B$ to anything to the left of $B$.

Recall that, in order for preferences to be single–peaked, the alternatives have to be lined up left to right ( in such a fashion that every voter has a single local maximum ). So the median alternative is the alternative that has half or fewer of the voters' preferred alternatives to its left, and half or fewer to the right. Alternatively : renumber the voters in order of their most preferred policy, so that voter #1 has the leftmost preferred alternative, voter #2 the second leftmost and so on. Then if there are 99 voters ( with voter #99 having the rightmost preferred alternative ), the preferred alternative of voter #50 will defeat any other in a vote. If the proposed alternative is to the right of her favourite, then voters #1 through 50 all prefer voter #50's preferred alternative ; if the proposal is to the left, then voters #50 through 99 will prefer voter #50's preferred alternative. In either case voter #50's preferred alternative gets at least 50 of the 99 votes.

Now there are many ways of lining up the alternatives. In figure 1, they were lined up in alphabetical order. But there is no good reason why they should be lined up in that order. The median voter theorem says that we can line up the alternatives in any order which we like. If there is **some** way of lining them up, such that **everybody** has single–peaked preferences, then there will be a Condorcet winner, namely the median of the preferred alternatives.

But with five (or more) alternatives, there are many ways of lining them up (120 of them, to be specific). If any one way of lining them up leads to all voter' preference profiles being single–peaked, then the theorem applies. So figure 2 does not prove that there is no Condorcet winner for that example : it just shows that preferences are not single–peaked when alternatives are lined up in alphabetical order. [In fact, alternative $D$ does happen to be a Condorcet winner in this example.]

It is a useful exercise to check the 3–person, 3–voter "standard" cycling example presented at the beginning of this note. With three alternatives, there are 6 ways of lining them up left to right ($ABC$, $ACB$, $BAC$, $BCA$, $CAB$ and $CBA$). You can check that, in this example, **none** of the six ways of lining up the alternatives gives rise to everyone having single–peaked preferences. Each way of lining them up leads to at least one voter having a graph of preferences which is not single–peaked.

How do we know that none of the six ways of lining up the alternatives in the standard cycling example will yield single–peaked preferences for everyone? Because, if there were just one way of lining up the alternatives such that everyone had single–peaked preferences, then the median voter theorem says that there is a Condorcet winner. And we already know that there is no Condorcet winner in this standard example of cycling.

## Making Economic Choices

Suppose that voters are choosing how much to spend on some public good. The public good must be paid for by taxes. The more we spend on the public good, the higher taxes will be. Each voter cares about how much of the public good is provided, but also about how much income she has left after taxes. If a voter expects to pay 3% of the cost of the public good, then she will have $Y - (0.03)Z$ left to spend on private goods, if a public good expenditure level of $Z$ is chosen — where $Y$ is the voter's before–tax income. More public expenditure raises $Z$, but lowers the amount of after–tax income that the voter has. The voter's consumption possibilities can be depicted in a graph, with public expenditure on the horizontal axis, and after–tax income on the vertical. Raising public expenditure will move her consumption down and to the right : there is more public good provided, but less income left after taxes. In fact, if the person expects that she will bear 3% of the taxes, then the consumption possibilities she has lie on a line with slope $-0.03$ — each dollar increase in public expenditure lowers her after–tax income by 3 cents.

There are an infinite number of levels of public expenditure to rank. But if the person has well–behaved preferences, then there is a unique alternative that she likes best. That is the policy which gives rise to the most–preferred consumption bundle on the graph. That's the point where her indifference curve is tangent to the line $X = Y - (0.03)Z$, if $X$ is her after–tax income. Now, as long as her preferences are convex, then moving right, down the curve $X = Y - (0.03)Z$ from her most–preferred outcome will move her to lower and lower indifference curves. So will moving left. Figure 3 illustrates : her most–preferred policy in that diagram is a public expenditure level of 2500, for which her indifference curve is tangent to her consumption possibility line $X = Y - (0.03)Z$. ( In the figure, the person's preferences can be represented by the utility function $U = XZ^3/100000000$. ) Increasing public expenditure above 2500 moves her to lower and lower indifference curves, as expenditure gets higher and higher. Decreasing expenditure below 2500 also moves her to lower and lower indifference curves, as expenditure shrinks further.

If she had to vote between two public expenditure levels, $Z_1$ and $Z_2$, and if $Z_1 < Z_2 < 2500$, then she would vote for the alternative $Z_2$ : both alternatives involve too little public expenditure from her point of view, but $Z_2$ is closer to her preferred level $Z = 2500$.

Figure 4 graphs this person's actual utility, as a function of the level chosen of public expenditure. The tax rules, her income, and her preferences are all the same as in figure 3. ( Her income is 100, she pays 3 percent of the cost of the public expenditure, and her utility function is $U(X, Z) = XZ^3/100000000$. So what is graphed in this figure is $(100 - (0.03)Z)Z^3/100000000$. ) The utility, as graphed, is single–peaked. It increases with the level of public expenditure, up to her preferred level of expenditure, $Z = 2500$. Then it declines with public expenditure, as its level rises too far above her preferred level.

What made the example work? The person's opportunity set was a straight line. That will be the case as long as she expects to pay, through her taxes, some fixed fraction ( 3 percent in the example ) of the cost of the public sector. Her preferences were "ordinary" well–behaved preferences, with the usual shape of indifference curves. As long as her indifference curves exhibit the usual diminishing marginal rate of substitution, then she will have a single preferred level of public expenditure, and her utility from different alternatives ( given that she has to pay some of the cost of the public sector through her taxes ) is single–peaked.

Requiring that preferences be single–peaked for general choices — candidates for elected office, favourite movie star, location for a new subway — may be very restrictive and implausible. But when the choice involves public expenditure on some good, then preferences will be single–peaked, as long as people's indifference curves have the usual shape. So pairwise majority rule may very likely give rise to cycling problems if we are trying to rank movie stars. But if we are choosing a level of public expenditure, with the public expenditure financed by some tax, then preferences will be single–peaked. The median voter theorem then implies that there is a policy which is a Condorcet winner : the median of the different voters' most–preferred levels of public expenditures.

## A Sort of Exception to this Single–Peaked Property

Suppose that the public expenditure is spent on primary education, and the voters are all parents with school–age children. There are some peculiar features associated with the public finance of education. Everyone who lives in a city gets the right to send their children to public schools there. But parents also can send their children instead to private schools. In this case, they have to pay tuition to the private school – but they still have to pay their taxes. Whatever level of expenditure $Z$ is chosen by the voters, the parents have a choice. They can send their children to the public schools : that means their children consume a quality $Z$ of education, and the parents have $Y - sZ$ left to spend on themselves, where $Y$ is their before–tax income, and $s$ is the share they pay of the cost of the public school system. Or they can send their children to a private school, which has a quality of education $\hat{Z}$, and which costs $F$ dollars in tuition. In this case, the children get a quality of education of $\hat{Z}$, and the parents have $Y - sZ - F$ to spend on themselves, since they pay both public education taxes and the tuition fees in the private school. Clearly [ why "clearly"? ], the parents will not choose to send their children to a private school if $Z \geq \hat{Z}$. But if $Z < \hat{Z}$, then it might be the case that $U(Y - sZ - F, \hat{Z} > U(Y - sZ, Z)$ so that they choose to send their children to private schools.

So if parents have a very strong taste for education, and if the public school system is not providing a very high standard of education, what will they do? They will choose to send their children to a private school. There is some level of expenditure on education $\bar{Z}$,which is a threshhold level. If the public school system spends less than $\bar{Z}$, then the parents will send their children to a private school, because they view the quality of the public school system as inadequate. It will always be the case that $\bar{Z}$ is strictly **lower** than the parents' preferred level of public education spending : they won't pay the added cost of private schooling unless the level of quality of public education is a lot less than their preferred level.

Is these parents' utility from public education spending single–peaked in the level $Z$ of expenditure on public education? To the right of $\bar{Z}$, it is : they're consuming the public education by sending their children to public schools, so that everything looks like the example of figures 3 and 4 above. But if $Z < \bar{Z}$, then they are choosing to send their children to private schools. The only impact that the level of public expenditure $Z$ has on them is through the taxes they still pay. If they are not getting any benefits from public education, then they will want their taxes as low as possible.

So the possibility of switching from public to private schools ( and the fact that parents must choose one or the other for their child ) make preferences potentially "double–peaked". If utility is graphed as a function of $Z$ here, it first falls with $Z$ as $Z$ increases from 0 to $\bar{Z}$. Then, to the right of $\bar{Z}$, at levels of $Z$ high enough to induce the parents to send the children to public schools, the utility starts to increase with $Z$, up to the parents' preferred level, and then to decrease again at levels of $Z$ above this preferred amount. So the graph of the utility function has two peaks, at $Z = 0$ and at the parents' preferred level of public education.

Who is the Median Voter?

Ignoring again this peculiar "private school versus public school" example, the main point of this section is that each voter's preference over ( a single category of ) public expenditure is single–peaked. So if voters use pairwise majority rule to decide public expenditure, the median of the preferred levels of public expenditure will be chosen. If that median level is denoted $Z^{med}$, it defeats all other alternatives. If someone were to propose some alternative level $Z' > Z^{med}$, then it would be defeated : all voters with preferred levels of public expenditure of $Z^{med}$ or less would vote for $Z^{med}$ over $Z'$, for these people $Z^{med}$ and $Z'$ are both more than their preferred level, and they will vote for the least excessive of the two alternatives. Similarly, a proposal of $Z'' < Z^{med}$ will be defeated by a coalition of all voters whose preferred level of public expenditure is $Z^{med}$ or more.

Different voters will, in general, have different preferred levels of public expenditure. What determines a voter's preferred level of public expenditure — the point at which her indifference curve is tangent to the opportunity set in figure 3?

Voters will have different tastes for public expenditure. For example, voters with many school–age children ( or legislators representing districts in which families have many school–age children ) will tend to want more expenditure on public education than voters with no children.

I cannot say too much about tastes. But the voter's preferred level of expenditure is where an indifference curve is tangent to a "budget line". It is the level of public good provision she would demand if she could buy it ( at a price of 3 cents on the dollar, in my example ). So, if the public good is a **normal** good, then higher income people will tend to have higher quantities demanded at given prices. That argument suggests that higher income people will tend to prefer higher levels of public expenditure — at a given "price" (3 cents on the dollar, in the example) of education.

But prices are not given. The "price" of public expenditure is the share the person pays of the total cost of the public sector. That in turn is determined by the share she pays of the taxes used to finance the public sector. In general, unless we use head taxes, that share of taxes will go up with a person's income. For example, if the public expenditure is financed by a proportional income tax, a person's "price" of the public good is her income, as a fraction of the total income in the jurisdiction. If instead a property tax were used, a person's tax price would be her share in the total property tax base. Since higher income people tend to spend more on housing, in this case as well higher income people will pay a larger share of the cost of the public sector than lower income people.

So the effect of income on a voter's preferred level of public expenditure could go either way. Most public goods are normal goods, which means that quantity demanded goes up with income. But the share of the cost of most public goods tends to go up with income, which means higher–income people face a higher price ( in taxes they pay ) for most public goods. Which effect predominates depends on how much taxes go up with income, on the income elasticity of demand for the public good, and on the price elasticity of demand.

What that means is that the preferred level of public expenditure need not increase steadily with income, and it need not decrease steadily with income. Even if all voters had the same tastes, their preferred level of public expenditure might be a U–shaped function of income, or an upside–down–U–shaped function of income, or something even more complicated.

What does that imply about the median voter theorem? The median voter theorem still holds. There will be a level of public expenditure which defeats any other in a pairwise vote. That level is the median of voters' preferred levels of public expenditure.

But it is **not** necessarily true that the voter whose preferred level is the median of all the preferred levels — the "median voter" — is actually the voter of median income.

LAETR : If the preferred level is a **monotonic** function of income ( increasing **or** decreasing ) then the median voter is the voter with the median level of income. But if the preferred level of public expenditure is U–shaped ( or upside–down–U–shaped ) as a function of the voters' income, then the median voters are not the voters with the median level of income.