Preference Revelation : (c) Complications and Difficulties

The preference revelation mechanism described in the previous two sections make it a **dominant strategy** for a person to tell the truth about the value she places on a public good. That is, if she understands the mechanism, and if the government commits credibly to obeying its own rules — the level to provide of the public good, the method of dividing up the cost, the rules for the additional pivot tax — then the person's own self–interest is best served if she announces exactly what her true demand for the public good is. Any attempt to misrepresent her preferences cannot make her better off, and may make her worse off. The fact that this is a dominant strategy for the person means that telling the truth is the best strategy for her, regardless what other people choose to do when they decide how to report their own preferences.

In this section, some extensions of the mechanism are discussed briefly, and also some potential weaknesses.

Sharing the Cost

In the mechanism presented in the previous sections, each person's tax had two elements. The first quantity was her share of the cost of the public good, and the second was the pivot tax that she would have to pay if her answer altered the provided of the public good.

It was assumed that her share of the cost was simply a fraction $1/N$, where N is the total number of people. That is, it was assumed that the cost was divided equally among all the people. This assumption is unnecessary. Instead, it could be replaced with the following rule : for each person i, there is some share s_i of the cost of the public good that the person must pay. So (in the variable–quantity public good of section (b) above), if the total cost of a level Z of the public good were cZ, then person i would have to pay $s_i cZ$ in taxes to cover her share of the cost. The share s_i does **not** have to be the same for all people. All that is required is that the shares do pay for the public good

$$s_1 + s_2 + \cdots + s_N = 1$$

and that each person knows her own share.

The rule for determining the quantity of the public good is not changed. The quantity Z^* satisfies the condition

$$v_1(Z^*) + v_2(Z^*) + \cdots + v_N(Z^*) = c$$

as before.

The basic idea of the pivot tax is unchanged. Now the rule for the pivot tax is changed slightly. The quantity $\tilde{Z}_1$ which would be chosen if person 1 were left out of the process is now determined as

$$v_2(\tilde{Z}_1) + v_3(\tilde{Z}_1) + \cdots + v_N(\tilde{Z}_1) = (1 - s_1)c$$

So the right hand side of the equation has been modified ; but it still represents the unit cost of the public good, excluding person 1's share of the cost. If $Z^* > \tilde{Z}_1$, then the pivot tax paid by

person 1 would be

$$PT = \int_{\tilde{Z}_1}^{Z^*} [(1 - s_1)c - (v_2(Z) + v_3(Z) + \cdots + v_N(Z))]dZ$$

so that again, the only change is to replace $\frac{N-1}{N}c$ with $(1 - s_1)c$.

With this slightly modified mechanism, it will still be the case that telling the truth is a dominant strategy for person #1 (and for any other person, if the pivot tax is defined analogously). Of course, people would prefer that they be assessed a lower share s_i of the cost. But given that the shares have already been determined, the best a person can do, when surveyed, is to reveal her preferences truthfully.

General Equilibrium Preference Revelation

In the "variable quantity" mechanism (presented in the second part of this section), people are asked to state their entire inverse demand curve for the public good, some function $v_i(Z)$ expressing how much they are willing to pay for a little more of the public good, as a function of the quantity provided. Of course, if the demand curve sloped down, they could just as easily state their regular demand curve, the quantity of the public good which they would be willing to buy at different prices, were the good excludable and available for sale on private markets. The willingness–to–pay curve $v_i(Z)$ is just the inverse of this demand function.

However, ordinarily quantity demanded depends not only on price, but on other variables, such as the prices of other goods and services, and on the person's income. If the public good were a **normal good** then increases in a person's income would shift the whole demand curve to the right. So we would have to write the inverse demand curve as $v_i(Z, M_i)$ if the good were normal, where M_i was the person's disposable income, and where $\partial v_i/\partial M_i > 0$ because the good is normal.

So changes in a person's disposable income will shift her demand curve for a good, unless the income elasticity of demand for the good were exactly zero.

But with the mechanism used here, the person probably will be charged a pivot tax. The amount of the pivot tax depends on the person's own answers to the survey, but also on everyone else's answers to the survey. Paying a pivot tax will lower the person's disposable income. So the amount that the person has to pay in pivot tax should affect her demand curve, if the good is normal. The higher the pivot tax, the further the demand curve shifts down and to the left (if the good is normal).

But the person does not know everyone else's answers to the survey, at the time that she must figure out her own announced demand curve. So, if the good is normal, she must know the pivot tax to correctly figure out her own demand curve, and this is information she will not have.

This problem can be avoided only if changes in her disposable income have no effect on the location of her demand curve for the public good — in other words, if the income elasticity of her demand for the public good is zero.

If the income elasticity of demand is non–zero, then the mechanism cannot be so simple.

For that reason, the mechanisms presented previously in this section are described as "partial equilibrium" preference revelation mechanisms, in that they ignore income effects. If income effects are significant, then a somewhat different mechanism is needed, a "general equilibrium" mechanism.

Unfortunately, general equilibrium preference revelation mechanisms do not work as well as partial equilibrium equilibrium mechanisms. It turns out not to be possible to design a general equilibrium mechanism which *i* takes into account these income effects ; *ii* always guarantees that enough money will be raised to pay for the public good ; *iii* makes telling the truth a dominant strategy for all people.

It is possible to design general equilibrium preference revelation mechanisms which satisfy properties *i* and *ii* above, for which telling the truth is a **Nash equilibrium**. That is, telling the truth is best for me, if I know that everyone else is telling the truth. This is a much less powerful property than having telling the truth as a dominant strategy. Now people will only be willing to tell the truth if they think other people are doing the same thing. With a dominant strategy, as in the partial equilibrium preference revelation mechanism, I am best telling the truth, even if I suspect that other people may not understand the rules, and may not themselves reveal their own preferences. Not so with these general equilibrium mechanisms.

Where does the Money Go?

The (partial equilibrium) preference revelation mechanism was constructed so that the first part of the tax, the people's cost shares, exactly paid for the cost of the public good, no matter what quantity of the public good was chosen. The second part of the tax, the pivot tax, will be positive or zero for each person. It cannot be negative. So the government, if it uses this mechanism, will be guaranteed to collect at least enough money to pay for the public good, more than enough if the pivot tax actually collects positive revenue.

What should be done with the excess revenue, the money collected through the pivot tax?

It cannot be given back to people. If I knew that I was going to get a share of any pivot tax revenue, then I should take account of this possibility in making my decision. Every extra dollar I pay in pivot taxes will actually get me back some money. But taking into account this effect will change my decision–making. It adds another term into my calculation, and it is no longer going to be true that telling the truth is a dominant strategy.

This will also be the case if I get any share of the pivot tax revenue collected from other people. If I am person #2, my answer to the survey affects the pivot tax schedule faced by person #1. If I, person #2, am going to get some share of that revenue, then I should take into account how my answer to the survey affects this tax yield. Again, taking this effect into account may have only a small effect, but it will change my optimal strategy slightly, away from telling the exact truth.

A similar problem will arise if the government uses the pivot tax revenue to pay for some other category of public expenditure. As long as part of the pivot tax money is going to fund

some service which is useful to me, or going to reduce my income taxes, then I should take into account how my answer to the survey affects the total pivot tax revenue. And taking this effect into account will alter slightly my incentive to tell the truth.

So the mechanism will work perfectly only if the government can commit not to spend the pivot tax revenue on anything useful to the people being surveyed. They could just throw the money away. A better approach would be to commit to spend the money on other people. They could, for example, make a deal with another jurisdiction, that the other jurisdiction would get their pivot tax revenue, and they would get the other jurisdiction's. Once this deal is completed, I do not care about how my answers to the survey affect pivot tax revenue. The revenue will be spent in another jurisdiction. I will get some money from another jurisdiction's pivot tax, but my answer to the survey will have no affect on how much gets collected in the other jurisdiction. Similarly, if the Ontario government used a preference revelation mechanism to find benefits from subway projects in the Toronto area, they could earmark the pivot tax revenue for expenditure in Northern Ontario (and perhaps earmark pivot tax revenue from a survey of Northern Ontarians' benefits from highway expansion to be spent only in the Toronto area).

Collusion

The (partial equilibrium) preference revelation mechanism is immune to manipulation by individuals. That is, if I understand the mechanism, there is no way that I can do better by lying than by telling the truth.

Unfortunately, it is not immune too manipulation by **groups** of people.

As an example, suppose that there were three people, and an all–or–nothing (fixed quantity) public good was being considered. The public good has a total cost of \$1200, so that it will be built if and only if the sum of the announced valuations of the three people sum to a number which is greater than or equal to \$1200. Suppose that the true valuations of the three people are \$700, \$350, and \$350. If each person tells the truth, then the project will be built, since the sum of the valuations is \$1400. But person #1 is pivotal. In this case

$$v_2 + cv_3 = 600 < \frac{N-1)}{N}C = \frac{2}{3}1200 = 800$$

If everyone tells the truth, person #1 will be assessed a pivot tax of \$100 ($= \frac{N-1}{N}C - v_2 - v_3 = 800 - 350 - 350$). She is better off getting a project built than not : she finds the project is worth \$700 to her, and she has to pay only \$500 in taxes (her share of \$400, plus the pivot tax of \$100). So she would not want to lie.

But she can benefit from making a deal with person 2 to lie, if she has some idea of other people's true preferences. She could offer person #1 a small bribe to change her answer v_2 from \$350 to (say) \$500. If $v_2 = 500$, and if persons #1 and #3 tell the truth, then no person is pivotal : here $v_1 + v_2 = 1200$, $v_2 + v_3 = 850$, and $v_1 + v_3 = 1050$, so that $v_i + v_j > 800$ for any two people i

and j : the project would still be built even if any one person's share of the costs, and announced benefits, were left out.

So if person #2 changes his answer from \$350 to \$500, then no–one is pivotal, and no pivot tax is paid. This change saves person #1 \$100, since it frees her of the pivot tax liability. The change does not hurt person #2 (or person #3) : the project would have been built even if person #2 had told the truth, so having him exaggerate his benefits does not change anything. Thus if person #1 pays person #2 \$50 to change his answer from \$350 to \$500, then both person #1 and person #2 are better off than if person #2 had told the truth.

In this example, the lie did not affect the outcome. But it might. Once people realize that they can collude profitably, and that collusion involves lying by one or more of them, then the mechanism no longer works.

So, if people can easily negotiate with each other about how they answer the survey, the preference revelation mechanisms presented in this section will be susceptible to manipulation by groups.