Economics 230a, Fall 2016
Lecture Note 3: Public Choice

Public choice is an integral part of the field of public economics, because it is the mechanism through which government policies are made in cases where market failures or distributional concerns require government intervention in the economy. We have stressed that a lack of markets also can mean a lack of information, so it is useful to ask how well public choice mechanisms work to translate individual preferences into publicly chosen policies. In particular, we may ask how well public choice mechanisms fare in terms of (1) rationality and (2) efficiency.

Voting and Arrow's Theorem

To begin, suppose that there are three individuals, 1, 2, and 3, and that public decisions among three potential outcomes (for example, a level of public spending), *A*, *B*, and *C*, are made using majority voting. Suppose that individual preferences satisfy:

1: $A \succ B \succ C$
2: $C \succ A \succ B$
3: $B \succ C \succ A$

where the operator "$\succ$" means "is preferred to" (with indifference being defined by the operator $\sim$ and weak preference by $\succcurlyeq$). Each individual has well-defined preferences over the three outcomes, but the ranking based on majority voting is plagued by "cycling": *A* is chosen over *B*, *B* is chosen over *C*, and yet *C* is chosen over *A*. This is the famous Condorcet paradox, which suggests that majority voting is not rational in the sense that it does not satisfy the simple property of transitivity, which would imply that if *A* is preferred to *B* and *B* is preferred to *C*, then *A* must be preferred to *C*.

Can we do better? Arrow's Theorem addresses this question, starting with the following axioms:

1. Unrestricted domain, i.e., any individual preference ordering can occur.
2. The mechanism must satisfy the Pareto principle: if no individual prefers outcome *X* to outcome *Y*, then neither will the social choice mechanism.
3. Independence of Irrelevant Alternatives (IIA): If the social choice mechanism results in a preference for outcome *X* over outcome *Y*, and then a new option *Z* is introduced, the presence of *Z* cannot alter the social ranking of *X* and *Y*.
4. There can be no dictator: there is no individual for whom social outcomes are determined solely by that individual's preferences.
5. The social choice mechanism must be rational, in that either strict ranking or indifference must be transitive – a property *not* satisfied by majority voting in the above example.

The theorem states that *no* mechanism based on individual preference orderings satisfies all of these axioms, so we must consider which of the axioms we might do without.

One might stop and ask how Arrow's theorem squares with our usual notion of maximizing a social welfare function. It would seem that the ranking of outcomes based on a standard social welfare function satisfies all of these axioms; while this is true, a social welfare function is not based simply on individual *orderings* of outcomes, but on the strength of these orderings as well. As it incorporates considerably more information, there is no inconsistency with Arrow's Theorem. On the other hand, realistic social choice mechanisms typically do not go beyond the use of ordinal preferences, although the literature on mechanism design has developed approaches aimed at eliciting expressions of true preferences over outcomes with regard to public spending, commonly referred to as Groves mechanisms and based on the same approach as Vickrey's "second-price" auction. Groves mechanisms operate by requiring each individual to compensate others for the losses others suffer as a result of taking that individual's stated preferences into account in public policy decisions.

Single-Peaked Preferences and Public Spending

One possible route out of our dilemma is to exclude the first axiom given above from our list of necessary conditions. For example, suppose all individuals have single-peaked preferences over potential outcomes, meaning that the outcomes can be ordered in a single dimension in such a way that there is a single, global maximum for each voter, with movements in either direction associated with less and less preferred outcomes. The following figure illustrates single-peaked preferences for three individuals over outcomes, *A*, *B*, and *C*, with valuations on the vertical axis.

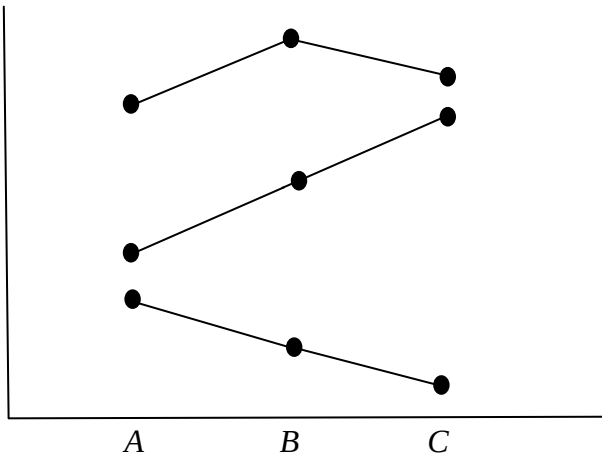

Note that for preferences to be single-peaked, there must be *some* ordering along the horizontal axis for which this property holds; it need not hold for *all* orderings of alternatives. Also note that the preferences listed above in presenting the Condorcet paradox are not single-peaked. There is no way of ordering the outcomes *A*, *B* and *C* in that example so that all three individuals have single-peaked preferences. Are single-peaked preferences a plausible restriction? Consider the case of individuals with constant budget shares voting over the level of public goods.

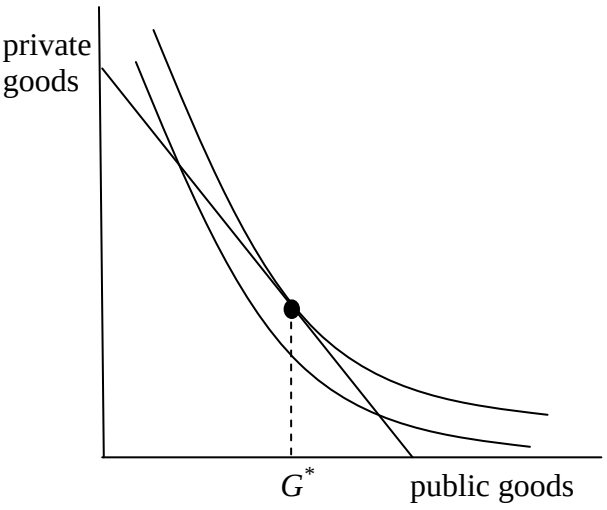

Suppose each individual has an endowment of private goods, and will be assessed taxes equal to a certain amount for each unit of a public good that is provided by the government. One might think of this assessment as coming through increases in tax rates, for example. The individual's preferred outcome is G^*, as shown in the figure at left, and movements in either direction, increasing or decreasing the level of *G*, will be associated with lower and lower levels of utility. Thus, the individual has single-peaked preferences over the level of public goods, and so will all other individuals, if they, too, face fixed assessments per unit of the public good.

Note that if preferences are single-peaked, then majority voting satisfies the remaining axioms, and the winner against all other alternatives will be the one preferred by the *median* voter – an outcome known as the median voter theorem.

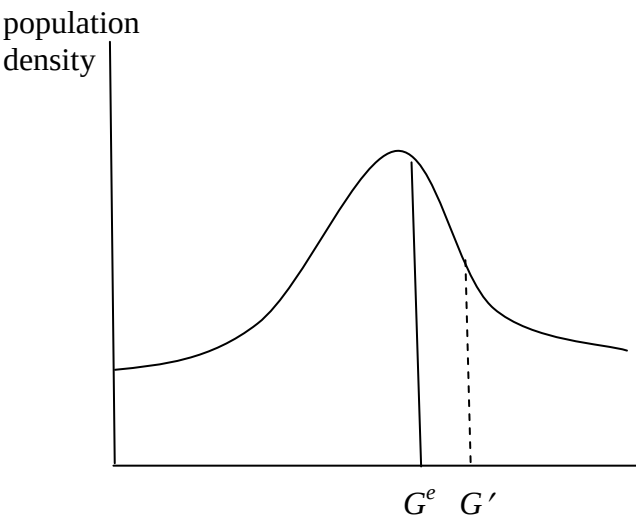

The diagram at left illustrates a possible distribution of preferred outcomes for a population, with the median voter's preferred outcome indicated by G^e. At this point, half of the population would prefer more and half less. Any option G' to the right would lose against G^e among the 50% of voters to the left of the median voter, plus some additional voters between G^e and G', and the same would be true of any option to the left of G^e. Thus, for any majority vote, or sequence of majority votes, including G^e as an option, G^e will be the winner.

As to whom the median voter might be, we often think of this individual as being the one with median income. This would follow if budget shares are equal, preferences are identical, and public goods are normal goods, for then higher income would lead to a higher preferred level of public spending. But if higher income is also associated with a higher budget share (as it would be under a progressive income tax) or if higher-income individuals have a weaker underlying taste for public goods, then whether higher income is associated with a preference for more public spending would be less clear.

Even though single-peaked preferences ensure that majority voting leads to stable, rational outcomes, there is no guarantee that these outcomes are efficient. This is a consequence of basing decisions only on *ordinal* preferences. At G^e, the median voter is at his or her preferred outcome for public spending – at a point of tangency in the previous diagram, where $p^h = MRS^h$: the individual's valuation of public goods in terms of private goods just equals the individual's marginal payment for public goods. For those who would have preferred a lower (higher) level of spending, $p^h > (<) MRS^h$. Assuming the budget shares add to the marginal cost of the public good, $\Sigma_h\, p^h = MRT$, we need that $\Sigma_h\, p^h = \Sigma_h\, MRS^h$ to achieve the Samuelson-rule outcome. But this will be true at G^e if and only if the sum of $p^h - MRS^h$ for those who would prefer less spending is just equal to the sum of $MRS^h - p^h$ for those who would prefer more spending. Absent such symmetry, we will get either too much or too little public spending. For example, if those who prefer more public spending do so mildly, but those who prefer less public spending have a strong preference in that direction, G^e will be above the efficient point.

Under the theoretical construct known as Lindahl pricing, each individual's budget share would be set equal to that individual's valuation of public spending at the socially optimal point, say G^L (i.e., $p^h = MRS^h(G^L)$), with those having a higher valuation of public spending (because of stronger preferences or higher income) being assessed a higher budget share, p^h. In this case, every individual's preferred outcome would be the same and equal to G^L. This level would win

not only by majority vote, but also by unanimous vote, and would coincide with the efficient level of public spending. Lindahl pricing is not achievable, because we don't observe preferences, but it serves as a useful benchmark in a number of contexts. For example, it corresponds to the notion of benefit taxation that one encounters in the literature (i.e., that individual tax burdens should be based on the benefits they receive from public sector spending.

Another problem with the median-voter outcome is that, once there is more than one dimension of choice – perhaps over two public goods rather than just one – preferences may no longer be single-peaked over bundles of outcomes. For example, suppose that there are two public goods, and that the possible outcomes for units of the public goods to provide are A = (0,0), B = (1,0), and C = (1,1). Suppose that individual 1 places no value on public spending and, having a positive budget share, has an ordering $A \succ B \succ C$. Suppose that for individual 2 the public goods are strong complements, so that the worst outcome is B. This individual might have an ordering $C \succ A \succ B$. Finally, suppose that individual 3 views the two goods as perfect substitutes, which could lead to the ordering $B \succ C \succ A$. This trio of orderings replicates that given above for the Condorcet paradox.

Can other mechanisms do better than majority voting? One important alternative, which arises in the context of fiscal federalism, is that individuals may "vote with their feet" by sorting among jurisdictions according to the public services the jurisdictions provide. Under certain assumptions, such sorting reveals individual preferences in much the same way that market decisions do, an outcome associated with the Tiebout hypothesis of efficient provision of local public goods, to be discussed later in the term. In Tiebout equilibrium, the taxes individuals pay in their desired locations are essentially Lindahl prices.

Further Topics

Within the broad field of public choice and political economy, a number of areas are particularly relevant to public finance; these are covered in the survey by Persson and Tabellini. One set of issues involves redistributive politics – questions where there are groups with different interests and decisions are made with regard to how resources should be divided among these groups. Among the divisions particularly relevant for public finance policy decisions are by income, age, immigrant status, and employment status. Another set of questions involves the process by which decisions are made: the role of lobbying, the form of government, etc.

One issue of particular current interest can be thought of as an example of the principal-agent problem, where government is the agent and voters are the principals. This involves the various budget rules and restrictions that are imposed on governments, such as spending and deficit limits. If individuals made policy choices directly, there would seem little argument for constraining decisions with such rules, for the rules are bound to limit flexibility in dealing with particular circumstances. For example, if public spending has a particularly high value in some realized state of nature, why should we restrict a government's ability to spend? There are a variety of answers to this question, generally relating to the role of government as agent. With different objectives (in part associated with a shorter time horizon due to the election cycle) and superior information, governments may pursue their own self-interest at the voters' expense and

not necessarily be held accountable. Thus, budget rules may sometimes be welfare improving, even if they limit the scope for desirable actions, as Besley and Smart show.

But an ongoing challenge is how to design budget rules that "work" in constraining governments from deviating from desirable policies and yet do not unduly constrain them from functioning reasonably well. Auerbach considers how to offset the incentives that politicians of either party have to overspend because they may lose control of the allocation of public resources in a future election. He finds that budget rules should take account of the future consequences of current actions, but that stronger weight should be given to shorter-run consequences, as future actions may reverse the effects of current legislation. Going beyond the paper, one might structure rules that impose a longer horizon on decisions that are more difficult to reverse, which may help explain why, at least in the United States, projections for public old-age pensions (Social Security), which play a role in potential policy reforms, have historically been made over very long horizons (typically 75 years).

We have many examples of constraining budget rules, including a series since the 1980s in the United States (beginning with the Gramm-Rudman-Hollings legislations setting deficit targets), the Stability and Growth Pact (SGP) in the Euro area (setting targets for the debt-GDP ratio and the deficit-GDP ratio) and the balanced-budget rules that most US states have imposed on themselves. One interesting contrast is between the US states and Europe. Whereas the SGP was adopted at the supernational level as a condition for membership in the Euro area, US state budget rules were adopted by the states themselves without coordination and with no role played by the US federal government. Many arguments for centrally imposed budget rules in Europe would seem to apply to the US as well – e.g., limiting cross-border fiscal shocks within a currency union; protecting members from having to bail another out – and yet the paths followed are different. One question that arises in this context is the relationship of governments at different levels. A problem with fiscal federalism is that lower levels of governments may face a soft budget constraint, perceiving correctly that some of their deficits will be covered by the central government, even if no formal arrangement to do so exists. The current situation in Europe with respect to Greece seems to be an illustration of this phenomenon. One possible solution to this problem is greater centralization of government functions, with restrictions on the scope of the activities of local governments. But there is a trade-off here, as local governments may be better adapted to respond to the heterogeneity of local preferences. See the paper by Besley and Coate.

One other question to ask is whether budget rules matter. One might be skeptical, given that there is often weak enforcement and the rules themselves are often repealed or replaced. Poterba's paper finds, at least for the US states, that budget rules do make a difference: states with stronger rules cut budgets more quickly in response to negative fiscal shocks. But these results don't necessarily carry over to other contexts, as state budget rules have existed for decades, and therefore are of much greater durability than existing national or supernational budget rules. Also, the effectiveness Poterba finds serves to illustrate a potential cost of budget rules: using the same source of variation across states, Clemens and Miran (*AEJ: Economic Policy*, 2012) confirm that states with more stringent budget rules have been less able to practice countercyclical fiscal policy.