More on the limits to redistribution

David Bevan
Department of Economics
Oxford University

Extended Abstract

There has been continued interest in the relation between income tax schedules and political processes for many years. One important strand of this literature follows directly from the work of Mirrlees (1971). Whereas he studied the tax schedule that would be chosen by a planner who was concerned to maximize social welfare, a number of subsequent contributions explored what would be the outcome to the same problem if a voting mechanism were invoked instead. Early examples of this alternative approach were Itsumi, 1974, Romer, 1975, Roberts, 1977, and Meltzer and Richard, 1981. These papers invariably restricted attention to a simpler tax schedule than the general non-linear schedule studied by Mirrlees, and focussed on a linear tax. For the most part, they also restricted attention to simple majority voting, though the idea was addressed that the voting population might not coincide with the tax-paying population.

Despite these simplifications, the results were sometimes complex, depending on what restrictions were imposed. The traditional approach to ensure that majority voting has a well-defined solution was to assume that preferences were single-peaked. It emerged from the preliminary studies that this could not be assumed in the tax problem, even if all the underlying functions were well-behaved. It then seemed difficult to obtain robust conclusions.

Subsequent work showed that a less demanding condition, variously formulated, but generalized as requiring single-crossing preferences, is sufficient to ensure that a stable voting equilibrium exists (Roberts 1977, Hemming and Keen, 1982 Gans and Smart, 1996). Given this condition – which seems fairly weak – there is, *inter alia*, a majority voting equilibrium, and the median voter is decisive.

Subsequent work has explored extensions to include modifications to the voting mechanism, inclusion of public provision of public goods, and of publicly provided private goods (Boadway and Marchand, 1995, Corneo and Gruner, 2000, Harms and Zink, 2003, Borck, 2005).

What this literature has neglected, in the main, is any comparison between the outcome driven by the assumed political process and that obtaining under a benevolent planner. The only exception to this has been the occasional comparison between the marginal tax rate chosen (for a linear tax) by the median voter and that which would be chosen by a planner with a Rawlsian objective of maximizing the welfare of the worst-off. It is useful, but hardly surprising, to learn that the marginal rate chosen by the median voter is lower than that chosen by the Rawlsian planner.

The objective of the present paper is to explore the relation between the voting equilibrium and the planning equilibrium more systematically. In particular, what sort of welfare function would imply the type of tax function that would emerge from a majority vote?

The paper examines these questions in respect of three types of tax design that might be chosen by a planner, or offered to a voter – a linear tax, a log-linear tax, or a non-linear tax of the type that would emerge from welfare-maximizing calculations. It also pays some attention to the way in which the government net revenue requirement may influence these relations.

One interesting result is that, for the types of assumption that commonly underlie this type of analysis, there may be a close correspondence between the welfare maximizing and the voting equilibria. Consider, for example, the case of a log-linear tax schedule, a lognormal skill distribution, and log preferences over consumption and leisure. Then a planner maximizing a Benthamite welfare function would choose exactly the same degree of progression as the median voter. This example can be extended to the case of any skill distribution that is symmetric in the logarithm of skill; obviously that includes the lognormal itself, but also a wider class, such as the two-parameter Champernowne distribution. What is more, it would hold approximately for a wide variety of empirically well-fitting distributions which depart from the lognormal in virtue of having (more or less symmetrically) thicker tails.

The paper then goes on to explore how the relation between the two types of schedule may vary as the assumptions are relaxed. These relaxations include a shift away from the assumption that individuals are completely selfish, so that the planner's preference for redistribution may be more or less rooted in the underlying preferences of individuals.

There are two common stories on the limits to redistribution. The first is that incentive constraints may restrict an egalitarian planner from achieving as complete a redistribution as the planner would desire. The second is that the political process – typically as modelled by majority voting – may render those considerations irrelevant, with the self-interest of the median voter determining the outcome. This paper explores the relation between these two types of restriction on redistribution.