Preface

Economists have paid some, but not much, attention to the charitable sector in today's economies. Also known as the voluntary or 'non-profit' sector, the activity of giving time and money for the benefit of other people, the environment and cultural assets has grown to such proportions that it effectively forces detailed scrutiny. Giving takes many forms, and this presents problems of defining just what the voluntary sector is. A major international research project – the Johns Hopkins Comparative Non-profit Sector Project – classifies voluntary organizations according to the following characteristics: formal activity in the sense of having rules, self-governing and independent of government (though not financially), acting primarily as a non-business, not distributing profits, and being voluntary in terms of donations of time or money or both (Kendall and Knapp, 1995). On this basis, the voluntary sector accounts for around 2 per cent of total employment in Italy and Japan, 4 per cent in the UK and France, and 6 per cent in the USA (Kendall, 1996). Across eight countries (UK, USA, Sweden, France, Germany, Hungary, Italy and Japan) the operating expenditures of the non-profit sector accounted for an average of 4.6 per cent of those countries' GNP in 1990, an absolute magnitude of $614 billion (Salamon et al., 1995). By any yardstick, the voluntary sector is huge, and it is growing.

The problem with prevailing measures of the size of the non-profit sector is that they do not measure the true 'social value' of the sector. Social value must somehow reflect the output of the sector. Yet the sum of donations and grants does not measure output; it measures input, that is, the cost of supplying charitable services. In fact it does not even fully measure inputs, since the value of volunteer time is not taken into account. Similarly, contributions to GNP are not measured by social value but, again, by the costs of supplying the services. This approach to GNP measurement is familiar: many economic activities are not bought and sold in the marketplace, so there are no 'revenues' to observe. In such circumstances it is commonplace to measure the contribution to GNP in terms of the costs of providing the service, as with public education and public health services. But the resulting measures are imperfect and potentially misleading. One aim of the current volume, then, is to pursue the idea of measuring the economic value of the charitable sector by looking at measures based on output, not input. Put another way, we ask the question: what are people willing to pay for the services provided by the charitable sector?

To our knowledge, the answers to this question, reported in Chapters 2–6 of this volume, represent the first attempt ever to measure the value of the charitable sector in terms of willingness to pay. In case this looks like a straw man, inventing a measure that has no particular rationale, willingness to pay is precisely the measure that is used to measure the output of the marketed sector of the economy. Willingness to pay (hereafter WTP) reflects individuals' preferences for a good or service, whether that be the contents of a supermarket trolley, the conservation of a historic building, wildlife preservation, or the provision of care for the aged. We therefore treat charitable services just like any other economic good. The interest lies in the fact that those services are not directly marketed, and hence we have to resort to techniques of 'non-market valuation' to find out the WTP for them.

Non-market valuation techniques involve discovering what people would be willing to pay if only there were a market. Broadly, two techniques are involved. The first looks for existing markets and asks if they embody in some way the value of the associated good or service we are interested in. Suppose the problem is the economic value of cleaning up air pollution. We do not buy and sell pollution, but we do buy and sell houses and we know that house prices reflect the neighbourhood amenities surrounding those houses, including the quality of the air. This 'revealed preference' approach, then, looks for an associated, or complementary, market and estimates WTP from observations in that market. The second technique is familiar to anyone who has ever carried out, or been the subject of, market research: we ask people for their WTP. Sophisticated questionnaires are constructed with the aim of either asking directly for WTP (what are you willing to pay?), or asking whether respondents are willing to pay a particular price (are you willing to pay X? – yes/no), or asking for individuals' rankings of alternative options where there is a link to the cost of providing the option. In the last case, WTP is inferred rather than stated directly. These contingent valuation and contingent ranking approaches have become very powerful in recent years and are widely used in environmental economics and health economics. This is the first time they have been applied to charitable services.

Because the detail of the valuation study is extensive, we have presented a summary of the results of the study in Chapter 5. Chapters 1–4 explain the analytical foundations for the study and the questionnaire results. Three sources of economic value are identified: first, the WTP of the general public to ensure that charitable services continue to be provided; second, the WTP of the beneficiaries of the services to maintain those services; third, the benefits to volunteers from the opportunities provided by charities.

A questionnaire-based approach is used to establish the WTP of the general public. But we treat the beneficiaries' WTP separately. Because there are so many different groups of beneficiaries, we cannot possibly survey them

all. We deliberately chose a group of users where we would anticipate serious potential problems of using our valuation techniques: the homeless. Our reasoning was that if our approaches worked in this complex and sensitive case, they would probably work in many others.

We make no apology for treating charitable services like other economic commodities, but we fully appreciate that some people will find the approach questionable. For example, the very reason we have a voluntary sector appears to be precisely because the market system does not provide those services. Why does the market system fail? It fails because many of the goods and services do not have apparent markets – recall the air pollution example. It also fails because market systems operate through prices and those prices may well exceed the ability to pay of the vulnerable groups who are the very targets of many charities. But great care needs to be exercised when using these failures of the marketplace to criticize the economic valuation approach.

First, we have already seen that there are often markets in the non-market service or benefit: the example of the housing market for air pollution can be extended. Historic houses and archaeological sites have implicit market value because people spend money travelling to see them. Their costs of travel provide a clue to finding their 'price'.

Second, even if markets are absent, it does not follow that people would not be willing to pay if there were a market. If we can find that WTP, then it may be possible to 'capture' at least part of it by establishing an institution that translates the WTP into actual payments. Charges for entry to parks, conservation areas, cathedrals and so on are all examples of 'capture'.

Third, charitable services are rarely provided by charities alone. Care of the elderly is provided by charities, by local government and by the private sector. There are potentially comparable markets out there. But this rationale has to be treated with caution. The users of charitable services are sensitive to who provides them, and there are distinct preferences for provision by one agent rather than another. Additionally, while the service may appear superficially the same, there are often marked qualitative differences. None the less, there is substitutability across a range of charitable services.

Fourth, as Chapter 5 concludes, charities compete for scarce resources. Money received from donations could have been used elsewhere. Money received from government grants could have been used to provide other public services. It is important therefore to ensure that charities are the best way of providing those services. But we cannot know that unless we know what the ratio of their output is for every pound or dollar they receive. It may be more efficient to provide a service through local government rather than via a charity, or vice versa. Cost-effectiveness indicators exist, but they have limited validity once it is recognized that outputs vary in the range of benefits

they provide and in the quality of service provided. We will not know more about comparative efficiency until we have some broader calculus of cost-effectiveness which, at the moment, does not exist on any widespread scale. We hope our work begins the process of achieving that.

Finally, the techniques we have used are familiar in the world of environmental economics. When environmental economics began there were concerns that the environment was being turned into 'commodity', and that this was illicit because the environment is 'beyond price', and somehow not to be brought within the measuring rod of money. Some critics still argue that. But the simple logic of opportunity cost – that whatever we do uses resources that could have produced some benefit elsewhere – remains. Anyone who argues that costs and benefits are irrelevant to social decision-making has to explain how cost and benefit can be ignored. All kinds of moral arguments can be invoked for doing so, but the fact that charitable services use resources to provide those services remains. This always means that those resources could have been used to provide some other service, which may just as easily serve some moral purpose. Morality is not irrelevant, but it must account adequately for opportunity cost.

Part I of the book is concerned with demonstrating the economic value of charities. Part II focuses on the next logical question: if people are willing to pay more than they actually pay for charities, how can this extra WTP, this surplus, be captured and turned into flows of income for charities?

Chapter 6 addresses the question of the role that government fiscal policy can play in stimulating the flow of resources to the charitable sector. It has long been argued that donations can be far more effective if they are 'tax-efficient', that is, more donations are generated if tax allowances on giving are provided. While it may seem obvious that giving tax incentives should increase donations, this is not borne out in practice. A higher rate of marginal tax in a context where there are tax allowances should increase giving because more tax is written off for each pound or dollar given. But higher tax rates also mean less after-tax income, so there is an income effect, which depresses giving. The two forces work in opposite directions. The empirical evidence reported in Chapter 6 indicates that the net effect of tax incentives on giving is positive for the UK: giving is higher in the presence of tax incentives than it would otherwise be. However, tax incentives do not provide a very large stimulus to giving. Indeed, the extra donations stimulated by fiscal incentives are not large enough to offset the associated loss of tax revenues to the Treasury. A possible reason for the limited impact of tax incentives is the relatively restrictive scope of mechanisms for tax-efficient giving in the UK.

In practice it is the efforts of fundraisers, rather than any tax incentives, which provide the greatest stimulus for philanthropic giving. Chapter 7 ex-

amines the relative efficacy of different fundraising methods, in particular those which involve direct face-to-face contact between fundraiser and potential donor versus those which rely on remote methods such as telephone, television, mailshots and so forth. The evidence shows that people are significantly more likely to give when they are approached face to face. However, the average size of donations received tends to be larger via remote fundraising methods. Overall, face-to-face methods present charities with the highest return to each fundraising approach. However, people's generosity diminishes markedly as they are approached an increasing number of times.

Chapter 8 considers the extent to which charities should target the population from which donations are sought. It is well known that many charities do this. But what exactly is the return to the charities from this targeting? Chapter 8 finds that targeting does indeed elicit more donations than 'scattergun' approaches, although some targeting variables are shown to be much more effective than others. But it also shows that larger charities have less to gain from targeting than smaller charities. Since targeting involves reducing the size of the population that is 'trawled' for donations, the effect of targeting for large charities is to lose out from the reduction of scale that comes with targeting and to gain from the targeted population. There is a trade-off. For small charities, the scale effect is less important than the targeting effect: they stand to gain most from targeting. Overall, Chapter 8 offers the elements of a theory of optimal fundraising.

Part III raises an issue for discussion: can our measure of the economic value of charities be used as a measure of social capital? There is now a substantial interest in the notion of 'social capital' – the 'glue' that holds society together and without which there is mistrust and social enmity that interferes with the smooth workings of the economy. Indicators of social capital are scarce. There are numerous negative indicators – measures of inequality, corruption and social breakdown such as crime, divorce and family disruption – but there are few positive indicators based on broad questionnaires about feelings of trust, or on measures of civil and political liberties. Yet the most obvious indicator of the extent to which human beings care for each other, the relationship of giving, has been ignored in the social-capital literature. We offer the observation that, in measuring the 'output' of the charitable sector, we are contributing to that measure of social capital. Chapter 9 raises the possibility of using the value of the charitable sector as at least a component part of the definition of social capital, a departure from the existing, and rapidly growing, literature on the subject.

Chapter 10 completes the book with some summary conclusions on the size of the charitable sector and on the mechanisms for capturing the surplus value we claim we have identified.