Journal of Health Economics 36 (2014) 164–173

Contents lists available at ScienceDirect

Journal of Health Economics

journal homepage: www.elsevier.com/locate/econbase

Aversion to health inequalities in healthcare prioritisation: A multicriteria optimisation perspective

Alec Morton *

Department of Management Science, Strathclyde Business School, University of Strathclyde, 16 Richmond Street, Glasgow G1 1XQ, United Kingdom

ARTICLE INFO

Article history:
Received 1 February 2012
Received in revised form 26 March 2014
Accepted 2 April 2014
Available online 13 April 2014

JEL classification:
I10
I14
I18

Keywords:
Health inequalities
Mathematical programming
Multicriteria decision analysis (MCDA)
Healthcare prioritisation
Healthcare resource allocation

ABSTRACT

In this paper we discuss the prioritisation of healthcare projects where there is a concern about health inequalities, but the decision maker is reluctant to make explicit quantitative value judgements and the data systems only allow the measurement of health at an aggregate level. Our analysis begins with a standard welfare economic model of healthcare resource allocation. We show how – under the assumption that the healthcare projects under consideration have a small impact on individual health – the problem can be reformulated as one of finding a particular subset of the class of efficient solutions to an implied multicriteria optimisation problem. Algorithms for finding such solutions are readily available, and we demonstrate our approach through a worked example of treatment for clinical depression.

1. Introduction

How best to take into account health inequalities is an ongoing issue in the theory and practice of health economics. In many developed countries, policy makers and the public seem to care about inequalities in population health (see Morton and Airoldi, 2009, for a discussion of the recent history of health inequality policy in the UK), and demonstrate this concern through both policy action and surveys of social values. This preference has been an important preoccupation for health economists over the last two decades or so (Williams and Cookson, 2000). However, it is still not clear how concerns about inequality should be captured in practical appraisal procedures, whether at the national level (for example, economic evaluation of technologies, development of national guidelines) or at the local level (prioritisation of spend by health authorities on the ground). A particular difficulty in connecting theory to practice is that health planners are often very reluctant to explicitly state parameters that reflect differences in importance of one section of the population rather than another: it is hard to imagine a Minister of Health announcing that a QALY accruing to a smoker was to be valued as 80% of a QALY for a nonsmoker, for example. Yet most forms of quantitative modelling seem to require such explicit tradeoff statements. A further difficulty in connecting theory to practice is that most health planners work with aggregate data and population averages, and do not have access to information about individual members of the population – indeed for many purposes, information about the distribution of health status in the entire population may simply not be collected, and planners have to rely on extrapolation from surveys or field studies. Yet welfare economic models construct their models of societal value by building upwards from an individual base (see Østerdal, 2005).

In this paper, we consider the question of how to prioritise in the face of incomplete information about values and aggregated information about population health through the lens of Multicriteria Optimisation (MCO). MCO deals with the formulation of and solution procedures for optimisation problems where there are multiple conflicting objective functions which cannot be completely traded off against each other. Solving a MCO involves identifying all solutions that are efficient in the sense that they are optimal with respect to some aggregate objective function within a family of possible functions, rather than optimising a

* Tel.: +44 (0) 141 548 3610.
E-mail address: alec.morton@strath.ac.uk

http://dx.doi.org/10.1016/j.jhealeco.2014.04.005
0167-6296/

single unique objective function. Where the family of objective function consists of all monotonically increasing functions, this MCO definition of efficiency collapses to the familiar concept of Pareto efficiency. Thus, MCO procedures side-step the problem of explicitly parameterising an objective function – the client for the analysis is presented with a number of "efficient", but possibly very different, solutions, between which they can choose directly. In some cases the identification of efficient solutions may be sufficient to arrive at a choice (e.g. if there a single, or single acceptable, efficient solution); but even in other cases, where multiple efficient solutions exist, there may be advantage in focussing discussions on concrete choices between alternatives, rather than on abstract and polarising questions of how much one values a health benefit to one sort of person rather than another sort of person.

This paper follows in a line of health economic papers that seek to apply mathematical programming approaches in health economics (Birch and Gafni, 1992, 1993; Johannesson and Weinstein, 1993; Stinnett and Paltiel, 1996; Earnshaw and Dennett, 2003; Anand, 2003; Epstein et al., 2007; Cleary et al., 2010). It has long been recognised that the ordering derived from the cost per QALY rule can be interpreted as a heuristic solution to a more general, but implicit, optimisation problem (Weinstein and Zeckhauser, 1973; Dantzig, 1998). These more general mathematical programming formulations model healthcare prioritisation as resource allocation subject to a fixed budget constraint, where the opportunity costs of choosing a particular healthcare investment are explicitly modelled. The mathematical programming framework can take account of issues such as indivisibilities, returns to scale, interactions between alternative investments, and the availability of recourse actions if investment decisions taken under uncertainty do not yield satisfactory results. The paper of Anand (2003) in particular is similar in spirit to the present work in that the author observes the healthcare resource allocation problem evokes conflicting values and considers how optimal solutions may differ depending on the choice of objective function. However, to the best of our knowledge, MCO has not yet been proposed in a health economics context.

Equity can be captured in various ways in the context of the mono-criterion models already proposed in the literature. For example, a social planner can impose through constraints that a certain amount of resources are devoted, or a certain amount of health is delivered, to some particular group (Stinnett and Paltiel, 1996). Such an approach may – arguably – be appropriate in certain situations (for example, if one population group deserves redress for some previous wrong, then a mathematical programme could be constrained so that they are compensated to the extent that they were previously wronged). But is has clear disadvantages, in general. For example, it may be simply impossible to find a solution that meets all constraints. Imposing constraints on lower levels of resource consumption in particular may have perverse effects, where required expenditure limits for one group may only be possible if they are provided with ineffective or indeed harmful treatment.

In this paper we seek to make a number of contributions:

- we introduce MCO concepts and relate these to the welfare economic theory of health as it relates to health inequality;
- we show that the commonly used but ad hoc approach of representing equity concerns through constraints recommends solutions that do not satisfy a multicriteria efficiency condition;
- we show through a worked example that MCO is workable technology for healthcare resource allocation, where project selection is modelled through continuous or discrete decision variables.

The structure of the paper is as follows. In Section 2 we outline a basic model of healthcare resource allocation. In Section 3 we present key MCO concepts and demonstrate how they relate to the model of the previous section. In Section 4 we present an example based on prioritising treatments for depression in England, and Section 5 concludes.

2. Our model

2.1. The (health-related) welfare economic frame

How should a social planner think about choice between alternative healthcare investments? A common approach in health economics is to assume that such investment decisions are only value-relevant insofar as they have influence on the health of a population captured through a Health-Related Social Welfare Function or HR-SWF. A substantial theoretic literature exists on the underpinning normative theory – see Culyer (1989), Wagstaff (1991), Bleichrodt (1997), Williams and Cookson (2000), Østerdal (2005) and Epstein et al. (2009) – as well as a growing body of empirical research (Dolan et al., 2005, 2008).

The core idea behind the HR-SWF can be presented as follows. Let $\mathcal{N}$={1, . . ., i, . . ., N} be the index set for the members of the population. A very general form for the HR-SWF is as follows:

$$\sum_{j \in \mathcal{N}} w_j u(h_j) \quad (1)$$

The variable $h_j \in \mathbb{R}_+$ represents a measure of the lifetime health for person j. The concave increasing function $u : \mathbb{R}_+ \longmapsto \mathbb{R}_+$ captures the idea that the more health someone has, the less valuable (to the social planner) a marginal increase is. The scaling factor w_j reflects that the health of certain people may be valued more than other people, because of certain characteristics of these people (for example, some of them may be smokers, or may have dependants). If $w_j = w \forall j$, then this HR-SWF is interpersonally anonymous in the sense that the same health benefit (for example, one QALY) is valued the same when one individual receives it as when another individual at same level of health does.

The HR-SWF presented in (1) is separable in the sense that it is possible to value a health gain to one person without knowing anything about what health gains accrue to other people. This is not as limiting as it might appear: many non-additive HR-SWFs (e.g. the CES function) can be transformed by a monotonically increasing transformation to an additive function, and are thus "strategically equivalent" to an additive function, in the sense that in any optimisation problem, the additive transformed version of the function can replace with original non-additive function without changing the optimal solution. There are however functions – in particular the Rawlsian maximin function – that cannot be be transformed thus, and there are situations where one may wish to model interdepencies between persons - for example in the case where there are "caring externalities" (Culyer, 1989) – and so this is not a completely vacuous assumption.

Such welfare models are indispensable for theoretic analyses, but have limitations for purposes of practical use in an appraisal or resource allocation context. A particularly practical difficulty with this model is that it seems to necessitate measuring the health of every individual in a population and planning based on that individual level data. This is unlikely to be possible. One way round this difficulty is to work with a simpler function. Where health improvements are marginal, this can be justified. In particular, for health improvements $\delta = (\delta_1, \ldots, \delta_j, \ldots, \delta_N)$ that are "small" in the sense that the first order Taylor series $u(h_j^0 + \delta) = u(h_j^0) + ((du(h_j^0))/dh)\delta_j$