Journal of Health Economics 58 (2018) 67–75

Contents lists available at ScienceDirect

Journal of Health Economics

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

Allocation rules for global donors

Alec Morton[a,*], Ashwin Arulselvan[a], Ranjeeta Thomas[b]

[a] Strathclyde Business School, University of Strathclyde, United Kingdom
[b] School of Public Health, Faculty of Medicine, Imperial College London, United Kingdom

ARTICLE INFO

Article history:
Received 11 May 2017
Received in revised form 1 February 2018
Accepted 2 February 2018
Available online 8 February 2018

Keywords:
Resource allocation
Development aid
Game theory
Global health
Bilevel programming
Stackelberg game

ABSTRACT

In recent years, donors such as the Bill and Melinda Gates Foundation have made an enormous contribution to the reduction of the global burden of disease. It has been argued that such donors should prioritise interventions based on their cost-effectiveness, that is to say, the ratio of costs to benefits. Against this, we argue that the donor should fund *not* the most cost-effective interventions, but rather interventions which are just cost-ineffective for the country, thus encouraging the country to contribute its own domestic resources to the fight against disease. We demonstrate that our proposed algorithm can be justified within the context of a model of the problem as a leader-follower game, in which a donor chooses to subsidise interventions which are implemented by a country. We argue that the decision rule we propose provides a basis for the allocation of aid money which is efficient, fair and sustainable.

1. Introduction

The last fifteen years have seen huge strides forward in humanity's ability to protect itself against disease. According to the United Nations (2015), improved prevention and treatment has saved around 37 million lives over this period and 6.2 million lives from malaria (most of the latter being those of children). 13.6 million people living with HIV are now on antiretroviral therapy: many of these people would be dead without treatment.

Although some of this progress is due to economic growth and technological progress, a considerable share is due to aid money from rich countries and philanthropic individuals. In particular, the Global Fund to Fight Aids Tuberculosis and Malaria has disbursed $27 billion to low and middle income countries to assist them in funding treatments for these three diseases since its founding in 2002 (The Global Fund, n.d.).

Of course, funding is not unlimited. Many rich countries are preoccupied with their own problems: slow growth, ageing populations, the hangover of the financial crisis, the threat of terrorism and the like. In this environment, overseas development aid can be a tough sell to a sceptical public. Private individuals are not constrained in the same way, but even the largest private fortunes are finite. This is reflected in the trends in development assistance for health (DAH). Between the 2000 and 2009 DAH rose at an annualised growth rate of 11.3% with an average annual increase of $290.4 mn (2015 US$) in the millennium development goals (MDG) focus areas and $98.6 mn (2015 US$) in the non-MDG focus areas. Between 2010 and 2015 however, the annualised growth rate fell to 1.2%, with DAH for HIV/AIDS and other focus areas (with the exception of maternal and newborn/child health) remaining flat or decreasing (Dielman et al., 2016).

How, then, should aid money be spent? Several commentators have advocated the use of cost-effectiveness analysis (Center for Global Development, 2013; Teerawattananon et al., 2013). In its simplest form, cost-effectiveness analysis involves prioritising investments in decreasing order of benefit to cost ratio, that is to say "value for money". In health economics in particular, the benefits are normally operationalised as some sort of health benefit: for example infections averted or life years or quality adjusted life years (QALYs) gained (for more comprehensive expositions, see Tan-Torres Edejer et al., 2003; Drummond et al., 2015; Neumann et al., 2016 and Wilkinson et al., 2016). While assessing these cost-effectiveness ratios for a given set of healthcare interventions can involve substantial clinical and epidemiological expertise, the underlying cost-effectiveness principle is familiar and intuitive.

In this note we take issue with the argument that donors should fund the most cost-effective interventions, with cost-effectiveness defined as total cost of implementation per unit of benefit, at least when the partner countries are middle-income or at the upper end of the lower-income scale. We maintain that a disadvantage of such

* Corresponding author at: Postal address: Room 720, Sir William Duncan Wing, University of Strathclyde Business School, 199 Cathedral Street, Glasgow, Scotland G4 0QU, United Kingdom.
E-mail address: alec.morton@strath.ac.uk (A. Morton).

https://doi.org/10.1016/j.jhealeco.2018.02.003
0167-6296/

rules is that they result in the crowding out of indigenous financing of interventions, and thus result in under-allocation of resources to healthcare. We propose an alternative rule for such countries whereby donors subsidise interventions which are "at the margin" from the point of view of the recipient country, that is to say in interventions which are only just cost-ineffective for that country.

To our knowledge this is the first analytic study of this issue, and accordingly there is little relevant literature to review. There is an empirical literature on the so-called "fungibility" of aid (e.g. Lu et al., 2010; Roodman, 2012; Van de Sijpe, 2013; Dykstra et al., 2015) but this is more concerned with empirical estimation of the extent to which external money displaces local financing. There is a classic technical literature on game theoretic analysis of the relationships between different players in the aid system (e.g. Svensson, 2000; Martens et al., 2002), but this differs from our study, which is much more operationally focussed on the derivation of decision rules for one party – the donor – in this relationship.

Our study can be seen as falling in a tradition of studies which seek to generalise the rules of cost-effectiveness analysis by showing how these rules can be derived within a formal optimisation model (e.g. Weinstein and Zeckhauser, 1973; Morton, 2014; van Baal et al., 2016, 2018; Morton et al., 2016). The optimisation framework we use here is bilevel programming, which has to our knowledge not hitherto been used in this context: specifically the model which we use can be considered as a bilevel knapsack model. There are quite a few bilevel knapsack models that are present in the operations research literature (Caprara et al., 2013). Most of them cannot be applied to model our application as these models involve the leader and the follower competing for a common resource or the follower is pessimistic with a conflicting objective. Dempe and Richter's (2000) model is perhaps the closest to the one that we are proposing. However, their model involves in the leader just determining the capacity of the follower's knapsack and it does not offer the flexibility in terms of individual project subsidies.

The overall structure of the paper is as follows. In Section 2 we present the intuition behind our rule and demonstrate it with respect to a realistic numerical example. In Section 3, to make our argument precise, we provide a formal model of the problem of donor aid allocation as a leader-follower (Stackelberg) game and show that our Section 2 algorithm can be understood as a heuristic version of the optimal solution to this game. In Section 4, we discuss the prospects for using the model we propose to guide resource allocation to countries. Section 5 concludes.

2. Concept and example

A popular approach to prioritisation in healthcare is to evaluate possible healthcare interventions (henceforth "projects") on the basis of their incremental cost-effectiveness ratio (ICER). This ICER is the ratio of the incremental costs to the incremental benefits relative to the current standard of care, with the latter measured in a metric such as QALYs. A common prescription is to invest only in those projects which meet some cost-effectiveness threshold. This idea has the merit of being both grounded in economic theory and also practically implementable: it informs decision making in many countries.

Suppose a donor (henceforth "D") wishes to interact with a country ("C") which makes decisions on this basis. How should D make decisions? What we propose in this paper is that D rather than selecting cost-effective projects on C's behalf, D should frame its role as deciding which projects to subsidise in order to make them cost-effective for C.

The intuition behind our proposed approach can be depicted visually in Fig. 1. In this figure the diagonal represents the line of marginal cost-effectiveness. The green squares above the diagonal represent the low CER projects which will be implemented by C (before subsidy) and the red circles below the diagonal represent high CER projects which will be rejected by C before subsidy. The decision of the donor is thus to choose which (red) projects to subsidise so that they are brought onto the diagonal line.

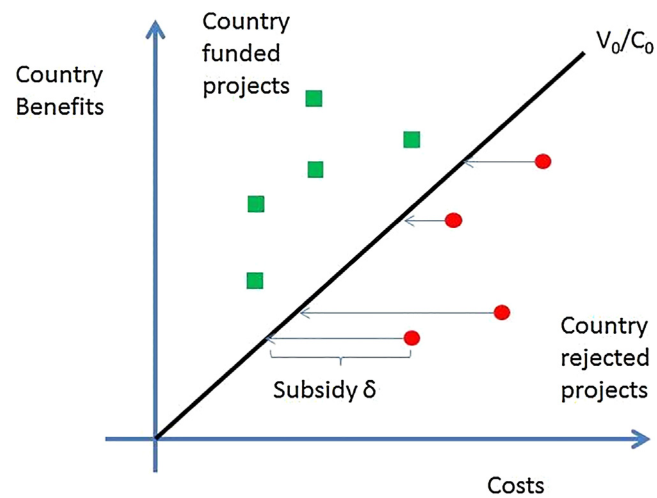

Fig. 1. Intuition behind the Donor's allocation rule.

This suggests an approach to prioritisation whereby D calculates its own Subsidy Effectiveness Ratio (SER) for projects by using the required subsidy, rather than the full cost of the project, as the numerator in the SER.

We demonstrate the approach numerically with the following example, based on earlier consultancy work. Consider a middle-income country which faces a range of projects in the areas of TB, malaria and HIV treatment and prevention, with costs and benefits (the latter measured in DALYs) and ICER as shown in Table 1. These projects may be considered to be targeted at separate populations and so are independent.

The first step is to order the projects in order of ICERs and identify those which have an ICER less than C's threshold, which we consider for the purposes of this example to be \$100 per DALY averted. These projects are shown in Table 2.

For the remaining projects, we now perform some additional calculations, shown in Table 3. The raw data for our calculations are the costs and benefits scores (columns 1 and 2) which we denote as c and b. To make the projects cost-effective for a country with a \$100 per DALY threshold, D will have to subsidise these projects by paying an amount δ equal to the ICER (column 3) of the project minus 100, all multiplied by the DALY benefits of the project (column 4) ($b(c/b-100)$). The balance of funds required to implement each project is met by C ($c-\delta$) and is shown in (column 5), which finds the subsidised projects (marginally) cost-effective. The donor then calculates a Subsidy Effectiveness Ratio (SER, δ/b) by comparing the benefits to be realised with this required subsidy (column 6). If for example, D has \$153 m, it may choose to implement the projects shown in bold on the table, which are identified by lining projects up in cost-effectiveness order and implementing projects successively in order of SER. We can verify (column 7) that this does indeed exhaust the budget.

In this example C spends approximately \$460 m on interventions which are already cost-effective for it, before D's financing begins to kick in. Had D spent money on these cost-effective interventions, it may well have crowded out financing from C. Instead, D's additional financing leverages substantial amounts of additional funding (\$215 m) into the healthcare system over and above what the country would have spent if left to its own devices. Thus,