

# Trading off health and financial protection benefits with multiobjective optimisation

Özlem Karsu\*, Alec Morton†

October 6, 2020

## Abstract

Countries which are introducing a system of Universal Health Coverage have to make a number of key tradeoffs, of which one is the tradeoff between the level of coverage and the degree to which patients are exposed to potentially catastrophic financial risk. In this paper, we first present a way in which decision makers might be supported to focus on in a particular part of the tradeoff curve and ultimately choose an efficient solution. We then introduce some multiobjective optimisation models for generating the tradeoff curves given data about potential treatment numbers, costs, and benefits. Using a dataset from Malawi, we demonstrate the approach and suggest a core index metric to make specific observations on the individual treatments. Moreover, as there has been some debate about the best way to measure financial exposure, we also investigate the extent to sensitivity of our results to the precise technical choice of financial exposure metric. Specifically, we consider two metrics, which are the total number of cases protected from catastrophic expenditure and a convex penalty function that penalizes out-of-pocket expenditures in an increasingly growing way, respectively.

*Keywords: Cost Effectiveness Analysis, economic analysis, financial protection, health benefits packages, Extended Cost Effectiveness analysis*

## 1 Introduction

The ideal healthcare system would pay for every effective medical treatment, for everyone in the covered population, entirely. Unfortunately in a world where budgets are limited, such generous coverage is possible only for the very wealthiest global citizens. Most existing healthcare systems, in both rich and poor countries, and everywhere inbetween, embody compromises - in terms of who is covered, what is covered, or what share of the cost is covered.

---

\*Department of Industrial Engineering, Bilkent University, Ankara, TURKEY. E-mail: ozlemkarsu@bilkent.edu.tr

†Management Science Department, University of Strathclyde Business School, Glasgow, UK. E-mail: alec.morton@strath.ac.uk

Recently, many national and global actors have fallen in line behind the UHC banner of “Universal Health Coverage” (UHC), a succinct label for the idea that everyone should have access to affordable and effective healthcare [WHO, 2010] [Chalkidou et al., 2016]. But what “affordable” means is highly situationally dependent. Moreover, even though universality takes the question of who is to be covered out of the picture, it gives little guidance as to how the balance should be struck between what is covered and what share of the cost is covered. Clearly when budgets are finite, if health services are provided free at the point of delivery, less can be provided than if the patients are required to contribute a share of the costs [Smith, 2013].

This tradeoff is recognised by Disease Control Priorities 3, the recent project of the World Bank and the Bill and Melinda Gates Foundation, which sought to provide analysis of healthcare “best buys” for Low and Middle-Income Countries (LMICs) [Jamison and DCP3, 2018]. The analysts working in this project have advocated a new form of cost-effectiveness analysis, so-called “Extended Cost-Effectiveness Analysis” or ECEA [Megiddo et al., 2016] [Megiddo et al., 2014] [Megiddo et al., 2018] [Morton and Lauer, 2017] [Nandi et al., 2016] [Verguet et al., 2016] [Verguet et al., 2015] [Verguet et al., 2013]. One of the defining features of ECEA is that financial protection is modelled explicitly alongside health benefits and expenditure. However, while ECEA clearly enriches the information base, decision makers may find it more taxing to move from the tables of numbers produced as part of a ECEA analysis to actual decisions about what to fund. (For comparison, it is relatively easy to “read off” the action implications of a cost/ Quality Adjusted Life Years (QALY) ranking.)

In this paper we offer a technical solution which can help decision makers make tradeoffs between health and financial protection. The sort of setting we have in mind is the determination of a “health benefits package”- an explicit list of all of the treatments to which a population is entitled, and at what copayment rates. The analytic and political challenges of determining a health benefits package are clearly an order of magnitude greater than those facing rich country Health Technology Assessment (HTA) agencies such as the NICE or the Scottish Medicines Consortium making yes/no decisions about particular technologies. However, for LMICs which have the ambition to rapidly scale up their healthcare provision in the face of a seriously constrained resource base, setting

explicit benefits packages may be attractive for several reasons [Glassman et al., 2017].

The tool which we propose in this paper is based on multiobjective optimisation. Mathematical optimisation methods have been used by many as a conceptual framework for understanding health systems [Morton, 2014] [Morton et al., 2018] [Morton et al., 2016b] [van Baal et al., 2016] [Smith and Yip, 2016] [van Baal et al., 2018]. The recent reports of the ISPOR Taskforce on Constrained Optimization point the way forward to using these methods more widely in operational analysis to support the development of more efficient, more patient-centred care [Crown et al., 2018] [Crown et al., 2017].

Multiobjective optimisation differs from traditional optimisation in that it seeks to generate not a single “best” solution, but a range of solutions which may be best depending on where the decision maker sets the tradeoff between competing objectives [Ehrgott, 2005]. This can be particularly useful in a portfolio choice setting such as the determination of a benefits package, where making choices “locally” about particular technologies using different criteria may result in global choices which are not logically consistent [Morton et al., 2016a] [Salo et al., 2011].

To illustrate how multiobjective optimisation would outperform a traditional approach, let us consider an example with four health care treatments that belong to two categories. Each treatment is evaluated with respect to its health benefit and financial exposure levels. In category 1 there are treatments A and B, and in category 2, C and D. One item has to be chosen from each category. The decision maker for category 1 subtracts the financial exposure score from the health benefits; the decision maker for category 2 subtracts twice the financial exposure score, resulting in the overall scores seen in Table 1:

Treatment	Health	Financial Exposure	Overall score
A	105	55	50
B	50	5	45
C	30	10	10
D	115	55	5

It is clear from the scores that one would choose A from Category 1 and C from Category 2. Let us take a look at the possible portfolios:

Portfolio	Health	Financial Exposure
A+C	135	65
A+D	215	110
B+C	80	15
B+D	165	60

As seen in Table 2 and Figure 1, portfolio (A,C) is dominated by portfolio (B,D) as the latter leads to more health and less financial exposure. Multiobjective optimisation is useful since one never chooses a dominated solution in the portfolio space because of local decisions guided by inconsistent values.

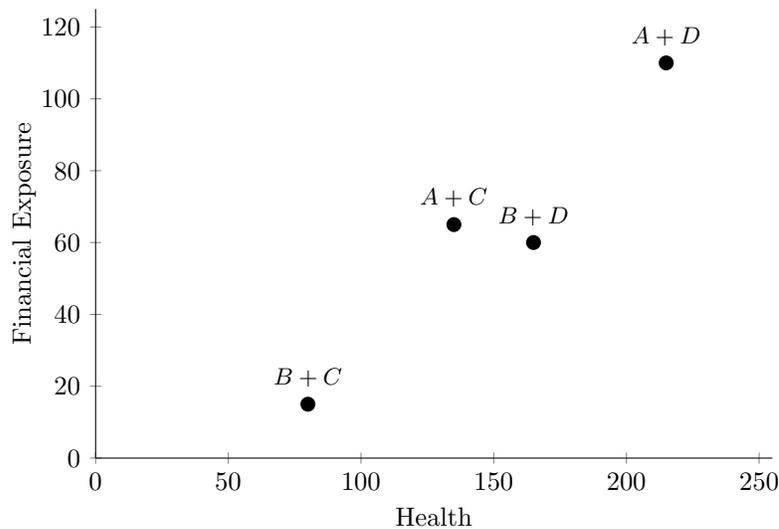


Figure 1: Bicriteria Analysis on Portfolios

This example demonstrates the complexity of the decision problem when multiple technologies are considered simultaneously. In such problems, one can use 2-dimensional ECEA graphs to visualize the policy maker's problem. However, it would still be challenging to evaluate all possible choice combinations (project portfolios), as the number of such portfolios would be prohibitively large. In real life settings with higher number of technologies to choose from, the complexity is

even more pronounced. Multiobjective optimization can be used to support decision making in such problems. Health benefits package design problem is a typical example where the policy makers have to compare many interventions simultaneously rather than focusing on one specific intervention. Indeed, any setting where decision makers have to make choices about multiple technologies simultaneously is a potential application for multiobjective optimisation. Further applications involve territorial resource allocation problem, in which a health authority with fixed funds prioritizes various health-related investments; and the specialty-level clinical planning, in which multiple improvement techniques (or multiple specialties) are prioritized so as to improve the healthcare system with respect to multiple criteria such as patient access and efficiency.

To our knowledge, the only paper in the health economics literature which has previously used multiobjective optimisation is [Morton, 2014], which uses the method to explore equity tradeoffs between different population subgroups who are differently disadvantaged. Morton’s paper uses both different model formulations and different solution algorithms from the present work.

In the following sections we will first discuss ways to support decision makers to choose their most preferred package given the tradeoff curve that shows how different options (health benefits packages) perform with respect to the two criteria of total health benefit and financial protection. We will then focus on the problem of finding this tradeoff curve, through a biobjective optimisation framework.

## 2 Interactive Methods

In this section, we suggest two approaches which may help a decision maker, who has to choose from a given set of solutions that are good to various degrees with respect to the two criteria of total health benefit and financial protection, narrow down which solution is the appropriate in the light of her objectives. Each such solution is a feasible portfolio of healthcare treatments (in the sense that the total cost of the treatments does not exceed the budget) and is also a Pareto solution, i.e., there is no other solution that improves it with respect to both criteria or with respect to one criterion without sacrificing from the other.

Our first approach is as follows. If the decision maker is prepared to assume an overall objective function of the form  $wH + (1 - w)FP$  (where  $H$  and  $FP$  denote the total health benefit and financial protection levels of a solution respectively), the situation becomes much simpler – even if she cannot specify the weight  $w$ . This is because the optimal solution must lie on the convex hull or envelope of the Pareto set. See Figure 2 for an example with 14 options.<sup>1</sup> Note that uncertainty in the weight does not necessarily translate into uncertainty in the decision: for example, whether the weight 0.75 or 0.85, the optimal choice is nevertheless option  $C$ .

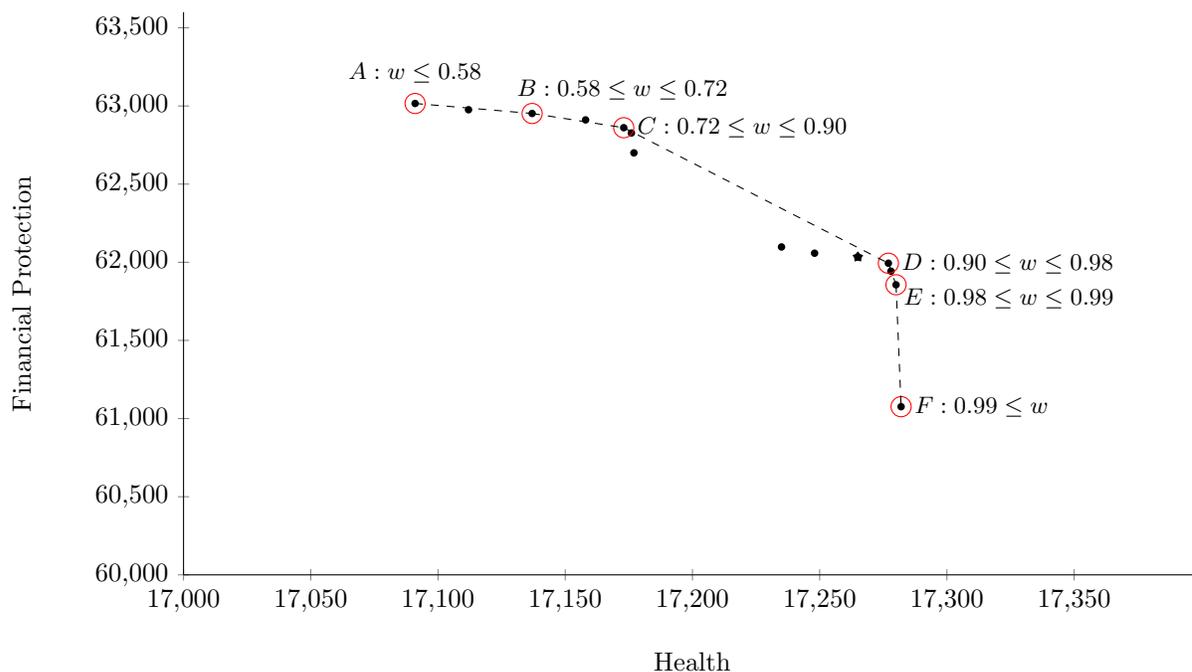


Figure 2: Tradeoff analysis

Our second approach uses the following idea, the history of which is discussed in [Karsu, 2013]. Assume that the decision maker’s objective function is  $f(H, FP)$  where  $f$  is some quasiconcave and increasing function, and suppose the decision maker has told us that she prefers  $a$  to  $b$  (see Figure 3). Then all points in the area  $bcod$  must be worse for the decision maker than  $b$ . To see why,

<sup>1</sup>This is the Pareto set of the example health benefits package design problem formulated with Malawian case study data with 4 coverage levels and a 90M budget, which will be explained in the following sections.

suppose that there was a point  $x$  contained in  $bcod$  which is superior to  $b$ . Then there would have to be an isoquant of  $f$  which separates both  $a$  and  $x$  from  $b$ . As is evident from the Figure such an isoquant could not be convex, which contradicts the assumed quasiconcavity of  $f$ .

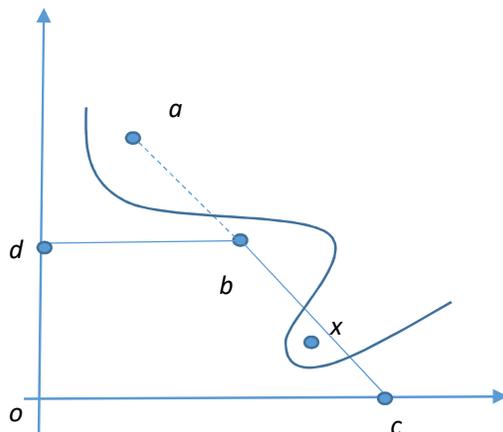
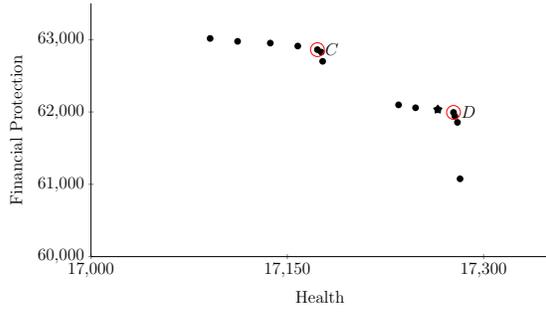
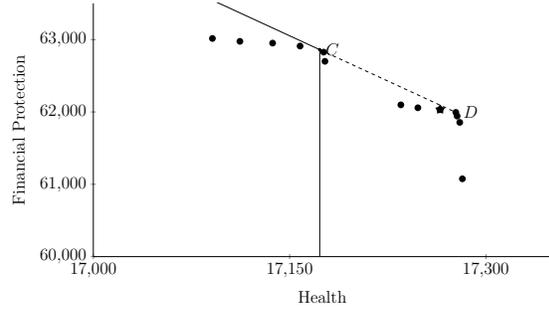


Figure 3: An example depicting the contradiction to quasiconcavity

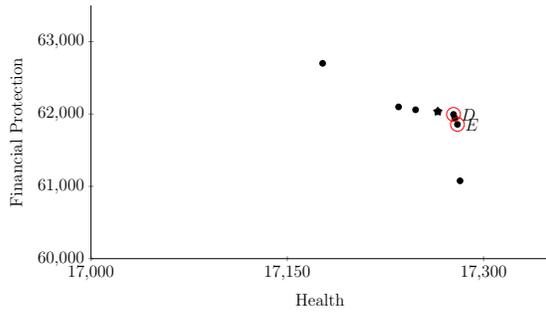
We do not provide full technical details of this idea which can be found elsewhere in greater generality than in the example studied in the current paper [Korhonen et al., 1984],[Karsu, 2013]. From our point of view, the important insight is that simple pairwise comparisons can eliminate whole regions of the Pareto set without the need to assume additivity of the decision maker's objective function. To get a sense of how this might work consider Figure 4. In the first step of this procedure, the decision maker is faced with a choice between  $C$  and  $D$ , and choosing  $D$  eliminates several points to the northwest of  $C$ . In the second step the decision maker chooses between  $D$  and  $E$  and her choice of  $D$  eliminates a further point to the southeast of  $E$ . The remaining five panels show how the remaining options are eliminated until a single circled point remains. The point of this example is to show that even if the Pareto set contains  $N$  points, it may not be necessary to explicitly examine each of those points to come to a decision.



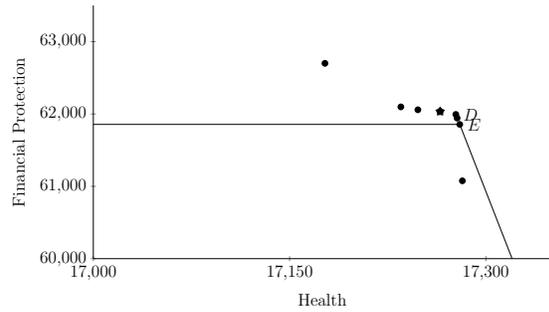
(a) Step 1: Choice



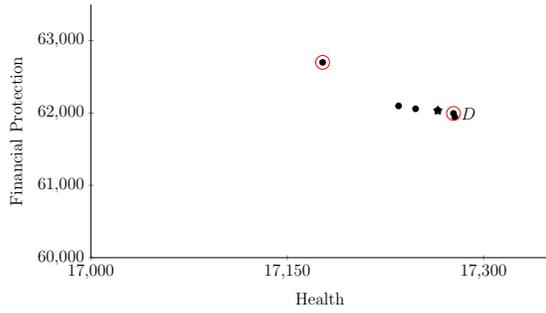
(b) Step 1: Elimination



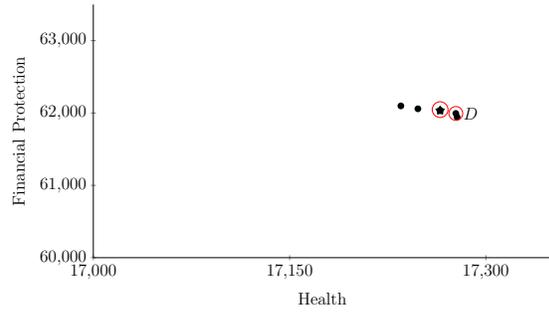
(c) Step 2: Choice



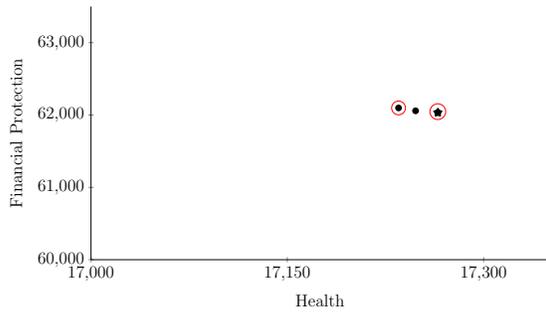
(d) Step 2: Elimination



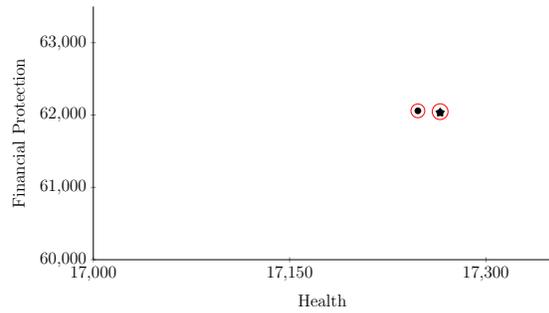
(e) Step 3



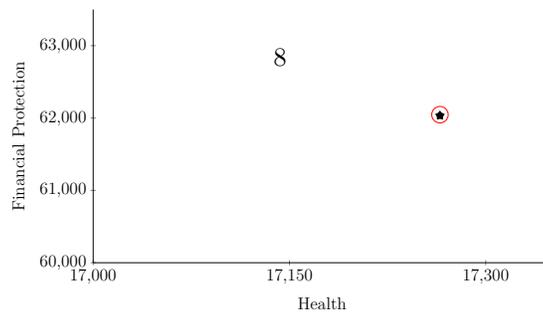
(f) Step 4



(g) Step 5



(h) Step 6



(i) Result of the interactive algorithm

A relevant design question is how to choose the pairwise questions to ask the decision maker to make the most out of her answers and determine her preferred option with the least number of questions possible. A good-rule-of thumb for choosing pairs for comparison would be checking the potential of the answer in reducing the number of options. As can be seen graphically from the figures, when choosing the pairs to ask, it is better to start with solutions that lie towards the center of the frontier as this would maximize the minimum number of alternatives that can be eliminated, i.e. it increases the potential to eliminate more alternatives. For example, in the first iteration, pair (C, D) is chosen because the minimum number of eliminations is the highest for this pair. If C is preferred then 4 options will be eliminated while if D is preferred 5 options will be eliminated. At least 4 alternatives are eliminated whatever the choice of the decision maker is.

Note that this rule prioritizes the solutions that lie on the convex hull of the set of (not-yet-eliminated) solutions. Moreover, it tends to choose closer solutions as they will reveal more information on the tradeoff between the two criteria.

Another question arises here on how to obtain the tradeoff curve, i.e. the alternative health benefits packages that are Pareto solutions. For this purpose, we propose using biobjective optimisation models with the objectives of maximising total health gain and minimising financial exposure (or equivalently, maximising financial protection). We discuss two such mathematical models that differ with respect to how financial exposure is quantified.

### 3 Proposed Models

Consider a setting where there are  $n$  treatments. Each treatment  $i$  is expected to generate a benefit value of  $b_i$  per person (measured in appropriate units) and consumes  $c_i$  units of resource per person. Let  $n_i$  be the total number of cases requiring treatment  $i$ . The total amount of budget available is denoted by  $B$ , which typically is not sufficient to cover all treatments to their full extent.

We assume that a central decision authority decides on the copayment rates, hence covering a portion ( $p_i$ ) of the costs and leaving the rest of the treatment costs to the patients.

We consider two concerns at the same time: The concern of maximising the total benefit achieved

by the package (measured in appropriate units such as total QALYs gained or Disability Adjusted Life Years (DALYS) averted) and minimising financial exposure. One can use different functions to assess financial exposure (see [Verguet et al., 2016]). In this paper we maximise cases of catastrophic financial burden averted (that is equivalent to minimising the number of cases with catastrophic financial burden). Any copayment that exceeds a certain threshold ( $T$ ) will be considered catastrophic. We will then discuss another approach based on minimising a convex penalty function.

Let  $p_i$  be the percentage of cost that will be covered by the package (i.e., 1-copayment rate) for treatment  $i$ . We will call  $p_i$  the *coverage rate*. Thus, if the coverage rate is 0.7 (70%), for an intervention costing \$15, \$10.50 ( $=\$15 \times 0.7$ ) will be paid by the government, and the patients will have to pay the balance of \$4.50 out of their own pocket. We assume that the decision maker chooses from  $m$  coverage rates as follows:  $0 (P_0) \leq P_1 \leq P_2 \leq \dots \leq P_{m-1}$ . Let  $M = \{1, 2, \dots, m-1\}$ . Let  $a_{ik}$ ,  $k \in M$  be a binary variable as follows:

$$a_{ik} = \begin{cases} 1, & \text{if } p_i = P_k \\ 0, & \text{otherwise} \end{cases}$$

Note that if all  $a_{ik}$  variables take a value of 0 for a treatment  $i$ , then the coverage rate is 0 for that treatment, which implies that  $P_0$  is chosen.

Let  $y_i$  be a binary variable as follows:

$$y_i = \begin{cases} 0, & \text{if treatment } i \text{ leads to catastrophic financial burden} \\ 1, & \text{otherwise} \end{cases}$$

$M$  is a sufficiently large number. The biobjective programming model we consider first is as follows. We call it the base model (B):

Model B

$$Max \sum_{i \in I} b_i n_i p_i$$

$$Max \sum_{i \in I} n_i y_i$$

$$s.t. \sum_{i \in I} c_i n_i p_i \leq B \quad (1)$$

$$p_i = \sum_{k \in M} a_{ik} P_k \quad \forall i \in I \quad (2)$$

$$\sum_{k \in M} a_{ik} \leq 1 \quad \forall i \in I \quad (3)$$

$$(1 - p_i)c_i - T \leq M(1 - y_i) \quad \forall i \in I \quad (4)$$

$$a_{ik} \text{ binary } \forall i \in I, \forall k \in M \quad (5)$$

$$y_i \text{ binary } \forall i \in I \quad (6)$$

Constraint 1 is the budget constraint. Constraint sets 2, 3 and 5 ensure that at most one coverage level is chosen for any treatment  $i$ . If none of the positive coverage rates ( $P_1$  to  $P_{m-1}$ ) is chosen, then no expenses are covered for that treatment, i.e. the copayment rate will be 1. For each treatment  $i$ , the model checks whether the patients are protected from catastrophic financial burden in constraint set 4. If the copayment  $((1 - p_i)c_i)$  exceeds the threshold  $T$ , the left hand-side of the constraint becomes positive, forcing  $y_i$  take a value of 0, meaning that the patients could not be protected from catastrophic financial burden. Otherwise the  $y_i$  variable will take a value of 1.

The two objective functions capture conflicting intuitions about how to prioritise health expenditures. The first objective function expresses the idea that the decision maker should seek to maximise the health gains which are attributable to public funding (hence the health benefits of the different treatments are scaled by  $p_i$ ). The second objective function captures the spirit of UHC: everyone should get the care that they need, but the cost burden on individuals should be managed so that as few people as possible are exposed to catastrophic financial burden.

One point needs to be clarified. In empirical work, the normal practice is to track catastrophic

costs which are actually incurred for healthcare. One of the limitations of such measurements is that it misses cases where people have the capacity to benefit from healthcare but do not have the means to purchase it, because of limited ability to pay. The interpretation of catastrophic financial burden in our second objective function is expenditure which is *required* because of some medical need. Whether this cost can or cannot be met is a secondary question: what is important is that there is disutility either because of untreated illness or because of financial distress.

It is possible to extend the model in a relatively straightforward fashion to the situation where one person can bear the cost of multiple diseases, and to the situation where there are multiple population quintiles. Rather than interrupt the flow of the narrative, we place these models in Appendix A.

As mentioned before, financial exposure could be assessed using various function forms. In the above models we have maximised the number of patients protected from financial catastrophe. The following model takes an alternative approach and minimises a financial exposure function  $f(\cdot)$  that is a convex function of the deviation from a threshold  $T$ . This allows us to go beyond simply designating a financial burden as being catastrophic or not and to incorporate different assumptions about “how catastrophic”.

$$\begin{aligned}
& \text{Max} \sum_{i \in I} b_i n_i p_i \\
& \text{Min} \sum_i n_i f(\text{Dev}_i^+) \\
& \text{s.t. Constraints 1, 2, 3, 5} \\
& (1 - p_i)c_i = T + \text{Dev}_i^+ - \text{Dev}_i^- \quad \forall i \in I \tag{7} \\
& \text{Dev}_i^+ \geq 0 \quad \forall i \in I \tag{8} \\
& \text{Dev}_i^- \geq 0 \quad \forall i \in I \tag{9}
\end{aligned}$$

Constraint set 7 is used to determine the difference between expenditures required for treatment  $i$  ( $(1 - p_i)c_i$ ) and the threshold  $T$ . Positive and negative deviations are denoted as  $\text{Dev}_i^+$  and  $\text{Dev}_i^-$ , respectively.

The above model is not necessarily linear due to the use of function  $f(\cdot)$ . To obtain a linear programming model, one can use piecewise linear convex functions that approximate (nonlinear) convex financial exposure functions as follows.

We first divide the range for possible values of  $Dev_i^+$  into intervals defined by levels. Let  $L_1, \dots, L_{int}$  be the deviation levels (see Figure 5 where  $int = 5$ ). In each interval between two consecutive levels, we assume linearity. Let  $w_1, \dots, w_{int}$  be the slopes of the corresponding intervals as seen in Figure 5. Since the function is convex, we have  $w_1 < w_2 < \dots < w_{int}$ . Given a set of levels and the corresponding slopes for the intervals, one can find  $f(x)$  for any deviation value  $x$  by detecting the interval  $x$  belongs to. For example, suppose that  $x$  belongs to interval 2 (i.e.  $L_1 < x < L_2$ ), then  $f(x) = w_1 L_1 + w_2(x - L_1)$ . In the mathematical model, we will define new decision variables that show the amount any deviation  $x$  spans in each interval ( $Dev_{ik}$  variables used below).

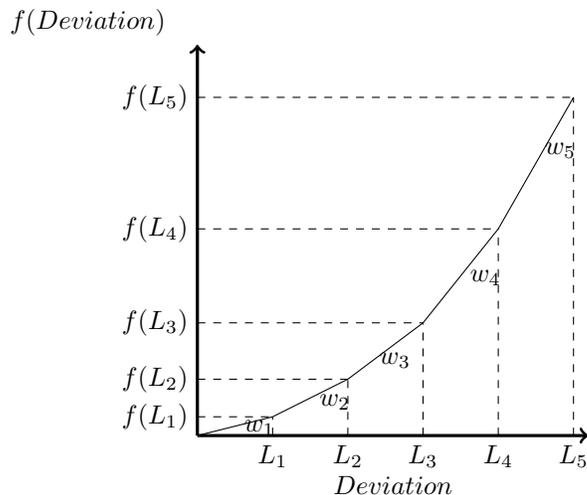


Figure 5: A Piecewise linear convex function with five intervals

The resulting linear programming model is as follows. We call it the convex financial exposure function model (C).

Model C

$$Max \sum_{i \in I} b_i n_i p_i$$

$$Min \sum_i n_i f_i$$

s.t. Constraints 1, 2, 3, 5, 7, 8, 9

$$Dev_i^+ = \sum_{k=1}^{int} Dev_{ik} \quad (10)$$

$$Dev_{i1} \leq L_1 \quad (11)$$

$$L_1 \delta_{i2} \leq Dev_{i1} \quad (12)$$

$$Dev_{ik} \leq (L_k - L_{k-1}) \delta_{ik} \quad \forall k = 2, \dots, int \quad (13)$$

$$(L_k - L_{k-1}) \delta_{ik+1} \leq Dev_{ik} \quad \forall k = 2, \dots, int - 1 \quad (14)$$

$$f_i = \sum_{j=1}^{int} w_j Dev_{ij} \quad (15)$$

$$\delta_{ik} \text{ binary } \forall i \in I, \forall k = 2, \dots, int \quad (16)$$

Constraint set 10 defines the  $Dev_{ik}$  values, which show the amount that  $Dev_i^+$  spans in each interval. Constraints 11-14 make sure that  $Dev_{ik}$  value for an interval  $k$  can not be positive unless the previous intervals are *filled up* first, using the binary variables  $\delta_{ik}$ s. That is, for example if  $Dev_i^+$  is in interval 2 ( $L_1 < Dev_i^+ < L_2$ ), then  $Dev_{i1} = L_1$  and  $Dev_{i2} = Dev_i^+ - L_1$ .  $\delta_{i2} = 1$  due to constraint 11 and the rest of the  $\delta_{ik}$  variables will have a value of 0. Note that there is no  $\delta_{i1}$  defined.

$f_i$  variables keep the per-person financial exposure level for treatments (That is,  $f_i = f(Dev_i^+)$ ) and are calculated using the  $Dev_{ik}$  variables and the corresponding interval slopes as explained above.

## 4 Numerical illustration

We illustrate our model with data relating to the Malawian health benefits package as described in [Ochalek et al., 2018]. The data table reproduced in Appendix B shows the list of treatments considered in this study (treatments that have very small number of cases or that have negative net DALYs averted are excluded).

To solve the biobjective model, we use the epsilon constraint approach [Ehrgott, 2005] using a stepsize of 1 for the total health benefit (since the number of cases values are scaled, this corresponds to an actual stepsize of 1000), to find the Pareto solutions. The approach is explained in Appendix C. The algorithm is coded in MATLAB and all mathematical models are solved in CPLEX 12.7. The algorithm is run on a personal computer with i5-4590 3.30 GHz processor and 8 GB RAM. Computation times are calculated in central processing unit (CPU) seconds.

We set  $T = 5$  and considered five coverage levels ( $M = \{1, 2, 3, 4\}$ ), which are 0, 0.25 ( $P_1$ ), 0.4 ( $P_2$ ) and 0.9 ( $P_3$ ) and 1 ( $P_4$ ) and report results for two different cardinalities of  $M$ . For the cases with three coverage levels ( $M = \{1, 2\}$ ), we used 0, 0.4 ( $P_1$ ) and 0.9 ( $P_2$ )<sup>2</sup>.

We considered three different scenarios for the total budget. Note that the total cost covering all treatments to the full extent would be around \$ 264.5 million. If all of them were covered to 90% that would require around \$ 238 million (M). We tried three levels for the budget parameter: 70 M (around 30 % of the total cost of covering all treatments to the highest coverage level), 90 M (38 %), 110 M (46%).

All of the computational times were under 5 seconds. The results for different budget levels (70M, 90M, 110M) and different number of coverage options (3 coverage level options and 5 coverage level options) are summarized in Figure 6.<sup>3</sup>

---

<sup>2</sup>These levels are set only for demonstration purposes. The framework can be used with any number of coverage levels.

<sup>3</sup>To ensure readability, we do not show one extreme solution of 70M,3 scenario which had  $7.44 \times 10^6$  DALYs averted and  $63.016 \times 10^6$  cases for which CFB is averted.

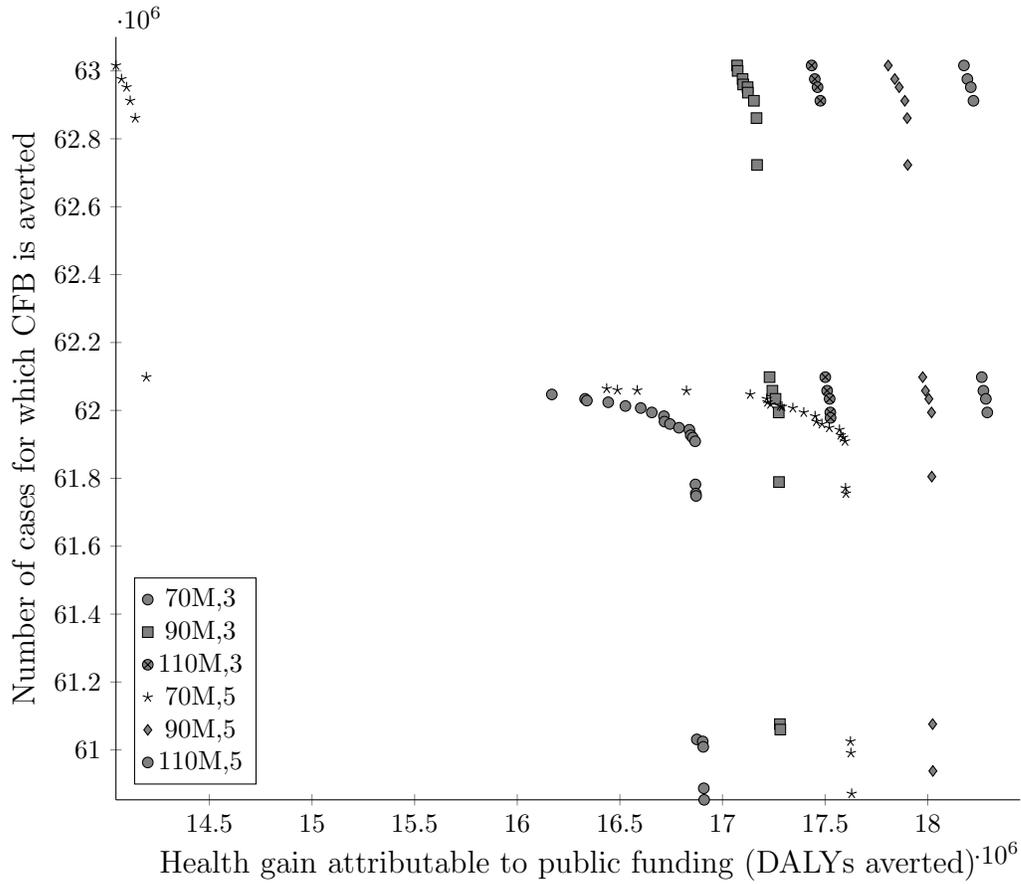


Figure 6: Pareto solutions for different budget levels and coverage level options. Each dot represents a nondominated (Pareto) point. (CFB: Catastrophic financial burden)

One can see that the solution sets for larger budget levels dominate the solution sets for smaller budgets, which is to be expected. The budget level also affects the number of solutions obtained. In general, we expect the number of solutions to be small for extreme (very high or very low) budget levels. Indeed, when the budget is sufficiently large to cover all treatments in the package fully, there is only a single solution with the best levels of both objectives. When the budget is very low, the number of feasible options, hence the number of nondominated solutions becomes small. At intermediate budget levels we expect the number of solutions to become relatively large as the

tradeoff comes into play (compared to high budget levels) and the number of alternative portfolios is larger compared to low budget levels. Note that the budget levels we chose for demonstration are all intermediate.

If, however, there are a small number of “large” projects and several “small” projects, which are included in some solutions while being left out at others, there may be sudden jumps at the frontier. See Appendix section D for a more detailed discussion.

Figure 6 also shows the effect of the number of coverage rate options ( $m$ ) on the results. One can see that for a fixed budget, increasing the number of possible coverage levels leads to better solutions in terms of both objectives. This implies that having flexibility with respect to the coverage options would increase the overall benefit of a universal health care package. On the other hand, using large numbers of possible copayment levels may not be convenient from a practical point of view. This is a design decision that should be made considering pros and cons of the options. Such analyses can help the policy makers see whether increasing the number of alternative copayment levels is justified.

Detailed analysis of solutions for the setting with 5 coverage levels and 90 M budget are seen below. Figure 7 shows the coverage levels of different treatments across the Pareto solutions. The treatments are ordered with respect to their core index (average coverage level across all Pareto solutions.) It is possible to detect the *must-do* treatments (2-7, 9,11-17,19-20,25-26,28,31,33-34,38-39,41-42) and the *do not do* treatments (18,23-24,36-37). Moreover, it is seen that treatments 1,29,35,46 are *fund at least partially* and 29,35,46,21,10,32,30,40 are *consider funding fully*.

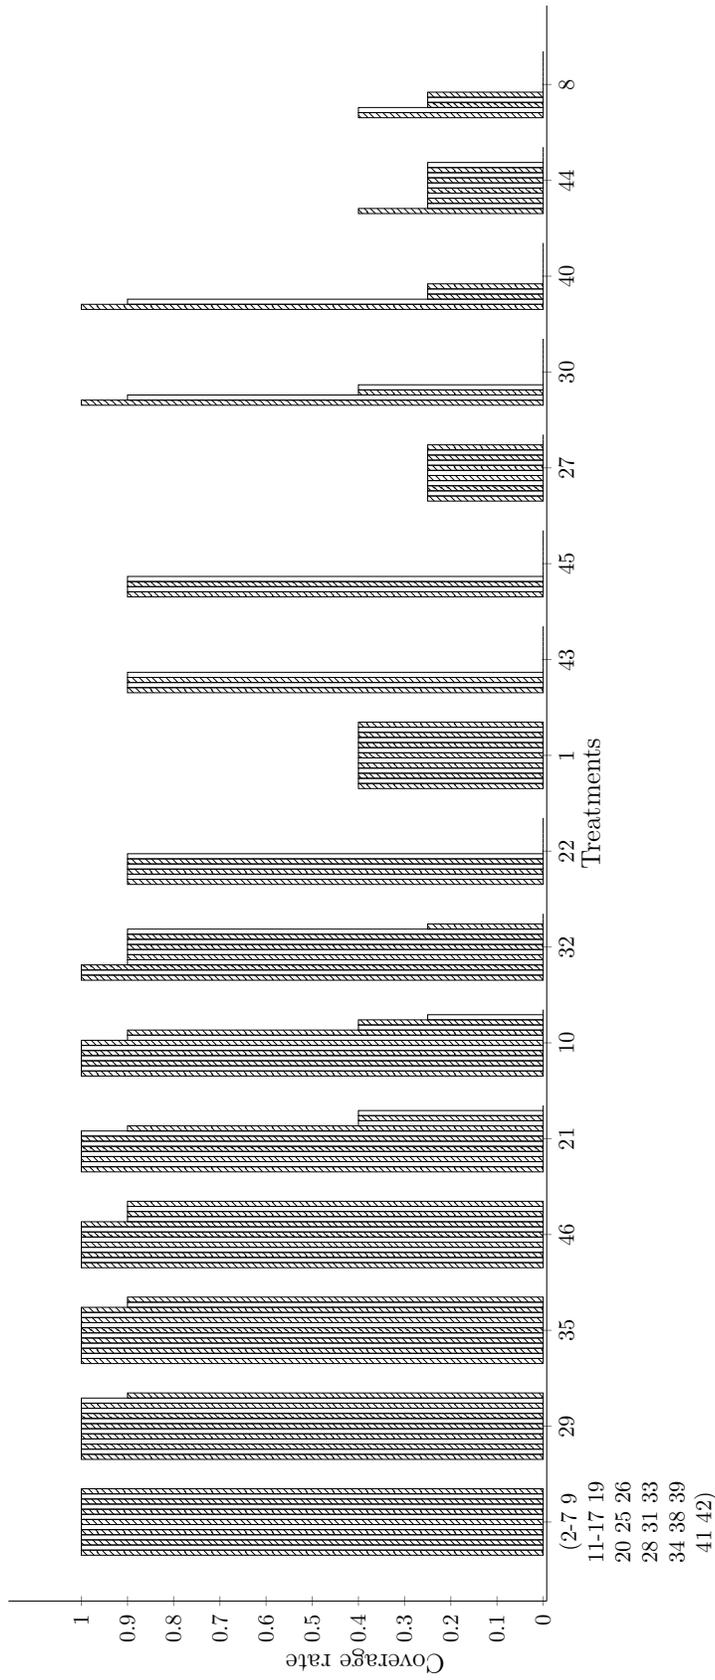


Figure 7: Coverage levels when model B is solved. Treatments (18,23-24,36-37) are not included in any of the solutions, hence they are omitted. Each bar shows the coverage rate in one Pareto solution. The bars only provide frequency information, i.e. the number of Pareto solutions in which a specific coverage rate is used for a treatment. The ordering of bars does not necessarily correspond to the same ordering of solutions across treatments.

For comparison, we re-did the analysis, this time using a convex financial exposure function. We again assumed that  $B= \$ 90$  million and there are five coverage levels 0 ( $P_0$ ), 0.25 ( $P_1$ ), 0.4 ( $P_2$ ), 0.9 ( $P_3$ ) and 1 ( $P_4$ ). We set  $T=5$ , i.e., the financial exposure function is a convex function of copayment amounts above 5.

10 levels are chosen for the piecewise linear financial exposure function, starting from 5 and increasing in increments of 5 units up to 50 (note that the maximum unit cost is 47.78). Figure 8 shows the function used and Figure 9 shows the results.

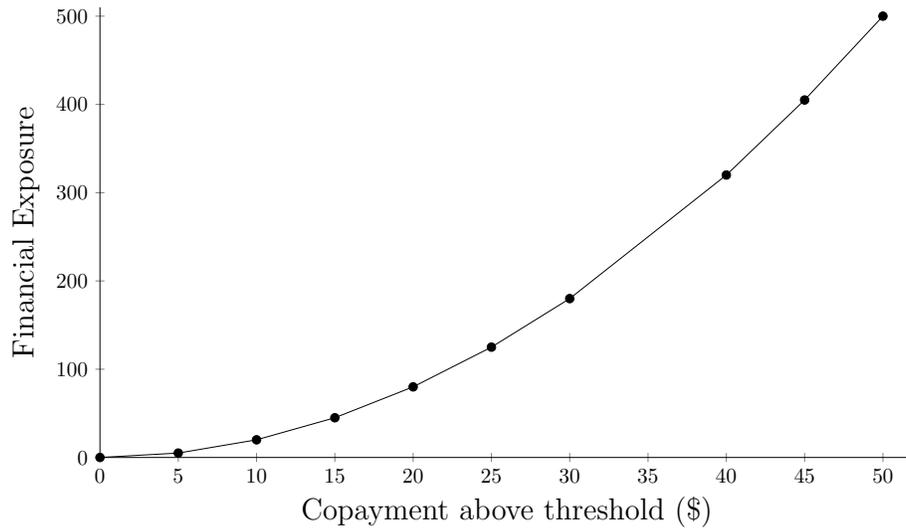


Figure 8: Financial exposure function used in the analysis

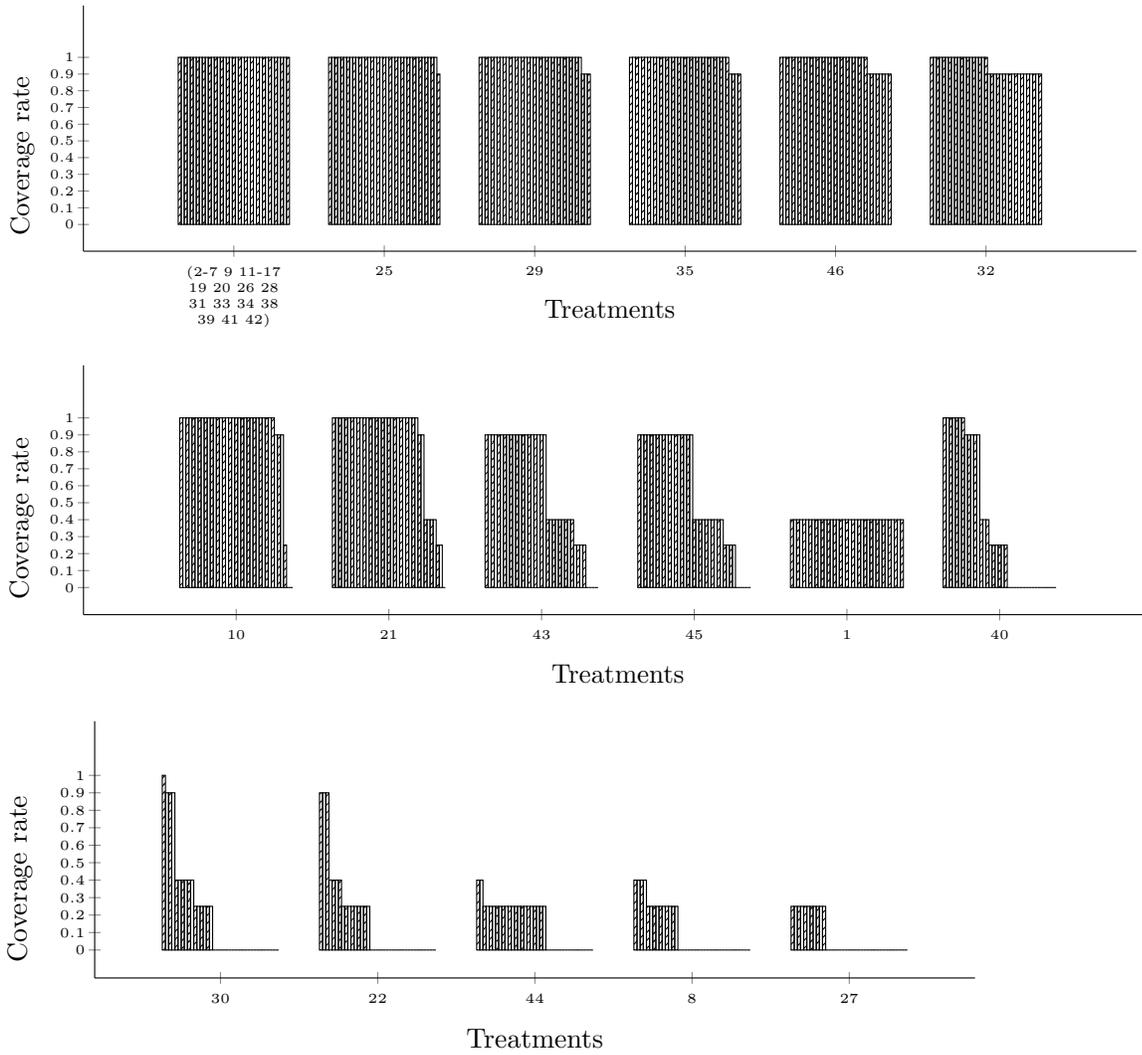


Figure 9: Coverage levels when model C is solved. Treatments (18,23-24,36-37) are not included in any of the solutions, hence they are omitted. Each bar shows the coverage rate in one Pareto solution. The bars only provide frequency information, i.e. the number of Pareto solutions in which a specific coverage rate is used for a treatment. The ordering of bars does not necessarily correspond to the same ordering of solutions across treatments.

In general, we observe that the solutions obtained using the two methods are very similar. However, the second method returns more solutions as it allows alternative solutions to be obtained by changing the deviation amounts above the threshold. In the first approach there is no distinction

between copayment levels that result in costs at the same side of threshold, they all contribute to the exposure function by the number of cases. However, how much one deviates from the threshold plays an important role in the second one. For example for treatment 43, one either uses a coverage level of 0 (in which case, the patient will have catastrophic financial burden) or 0.9 (in which case the patient will not have catastrophic financial burden) in the first method: 0.25 and 0.4 levels are not sufficient to save the individuals from catastrophic expenses. However, when the second method (convex penalty function) is used, such levels also matter since they lead to different amounts of deviation from the threshold. We obtain alternatives where 40% of the expenses are covered for treatment 43 when the second exposure function is used. Similar observations can be made for another high-unit-cost treatment, treatment 45.

To further the case that models B and C are closely equivalent, we ran both models for a number of different problem instances, obtained the ranks of treatments with respect to their core index and calculated the correlation between the rank orders obtained in both models. For this purpose we generated different scenarios with 25, 50, 75 and 100 treatments. For each level of  $n$  we created 5 problem instances, in which  $n_i$ ,  $c_i$ ,  $b_i$  are randomly generated between ranges [2, 16000], [1,50] and [1,60] (These ranges are similar to the ranges in Malawi data). We set the budget as 50% of the total money that would be required if all treatments are funded fully and set  $T=25$ . The average and minimum correlation values for different sized instances are seen in Table 3. The results indicate that the solutions are not sensitive to the metric used to measure financial exposure.

Table 3: Rank order correlations

$n$	Avg. Correlation	Min. Correlation
25	0.894	0.829
50	0.958	0.945
75	0.952	0.921
100	0.961	0.932

The solution times were negligible for smaller instances (under five seconds for  $n = 25$  and under five minutes for  $n = 50$ , using a stepsize of 1.). As the problem size increases the solution time increases for both models but more so in model C. This is because model C typically has more Pareto solutions than model B and each single objective model solved to obtain these solutions

is larger with more decision variables and constraints compared to the ones in model B. In larger instances, we increased the stepsize for model C to obtain the Pareto solutions in reasonable time. For  $n = 75$ , we increased the stepsize to 1000 for model C and the average solution times for were 1000 seconds and 3500 seconds for models B and C, respectively. For  $n = 100$ , we increased the stepsize to 5000 for model C and the average solution times were 2640 seconds and 2765 seconds for models B and C, respectively.

## 5 Conclusion

The purpose of this paper has been to study how decision makers might tradeoff health and financial protection benefits in the context of the construction of a health benefits package for a country which is moving to UHC. Our reason for using a multiobjective optimisation frame to study this problem is that making an explicit tradeoff between health and financial protection may be extremely difficult. Nevertheless, analysis can still provide assistance in helping decision makers avoid health benefits packages which are internally inconsistent. Our paper is a first step in the analysis of this question and allows plenty of scope for additional work.

Below, we discuss a number of extensions that could be handled within the multiobjective optimisation framework by modifying the mathematical models that we use.

Ministries of Health typically have as an objective the maximisation of population health when prioritising the funds which are under the control of the Ministry of Health. We interpret this as meaning that the Ministry wishes to maximise health benefits attributable to these funds and define the total health benefit objective accordingly. All out-of-pocket payments are accounted irrespective of whether they are below or above the threshold, when calculating the total benefit of the healthcare package attributable to public funding. Out-of-pocket payments below the threshold do not increase the second objective, which is related to the financial burden, as this objective focuses only on the catastrophic costs. It is, however, still possible to accommodate different views on how total health benefit or financial burden should be calculated by modifying the mathematical models, and changing the forms of the objective functions. For example, the Ministry of Health

may choose only to count the health benefits when they are entirely paid for public funds and no copayment is required; or the financial protection function could be a more smooth nonlinear function of the cost burden on patients. Hence the proposed biobjective optimization framework would still be relevant and useful.

Note that there may be interdependencies between treatments, making the costs, benefits or demand of treatments depend on the decision on other treatments. In such cases, the mathematical programming formulation could be modified to include these sorts of interactions, as long as they could be measured (the interested reader is referred to [Winston, 2004] for more discussion on mathematical modelling). Such interdependencies between interventions are addressed in the paper by [Hauck et al., 2019], where different ways of modeling these interdependencies and incorporating them into an optimization framework are proposed. In the health benefit package design problem that we consider, as long as measurement can be made, it is possible to incorporate these interactions into the optimization setting in similar ways to those discussed by [Hauck et al., 2019].

Further research could be performed in this domain to address various related questions. For example, it would be important to know how robust our models are to a range of approaches to modelling the financial protection benefits. Also, benefits may differ according to what delivery platforms and human resources are available [Morton et al., 2016b] [van Baal et al., 2018] [Hauck et al., 2019]. Furthermore, in a development setting, there may not be a single decision maker as a significant share of the funding may come from international donors [Morton et al., 2018]. Lastly, establishing a benefits package in an environment where there is copayments involves making tradeoffs between different population subgroups with different levels of ability to pay [Morton, 2014]. Equity concerns for patient groups with different abilities to pay can be incorporated into the biobjective optimisation framework by defining catastrophic expenses in relation to ability to pay, i.e. defining population group based thresholds, as discussed in Appendix A. Such a model would decide whether a person faces catastrophic burden based on the socioeconomic level of that person. Maximizing the number of people saved from such expenses would implicitly prioritize the ones with low thresholds. Compared to the base approach that does not take the ability to pay into account when defining catastrophic expenses, such an extension may change recommendations (the

set of Pareto solutions) in favor of lower income quintiles, however, we expect the main observations made on the tradeoff between total health benefit and total financial protection to still hold. We leave the integration of these complexities into our models as a challenge for further research.

## Acknowledgement

The authors are grateful to Dr Itamar Megiddo for his feedback on the initial version of manuscript.

## Appendix A Model extensions

### A.1 One person bears the costs of multiple diseases

In that case, in addition to the set of treatments  $I$ , one can generate a set of sets of treatments that have an intersection population (Let this be  $I_{ext}$ ). For each element  $s_l$  of this set, let  $m_l$  denote the number of people that are affected by all the diseases in set  $s_l$ . For example, consider a case where there are three treatments (diseases) and assume that there is a sub-population which are susceptible to diseases 1 and 2 together. Then  $I_{ext} = \{1, 2, 3, \{1, 2\}\}$ .  $m_1$  denotes the number of cases of disease 1 only,  $m_2$  denotes the number of cases of disease 2 only,  $m_3$  denotes the number of cases of disease 3 only and  $m_4$  is the number of cases where diseases 1 and 2 occur simultaneously. In such a case, constraint sets 4 and 6 are modified as follows:  $\sum_{i \in s_l} (1 - p_i)c_i - T \leq M(1 - y_l) \forall l \in I_{ext}$  and  $y_l$  binary  $\forall l \in I_{ext}$ . Of course, the second objective function should be as follows:  $\sum_{l \in I_{ext}} m_l y_l$ . We assume that the benefit function is an additive function, the total benefit that a patient with multiple diseases will get is the sum of the benefits of the treatments that are funded, hence the first objective function stays the same.

### A.2 Data supports analysis at population subgroup level

We suppose that (as in various ECEA applications) the relevant subgroups are socioeconomic quintiles. This would require defining quintile-specific parameters and decision variables as follows.

Let  $n_{iq}$  denote the number of cases requiring treatment  $i$  in income quintile  $q$  and  $Q$  be the set of income quintiles. One could define income specific thresholds for quintiles:  $T_q$  for all  $q \in Q$ . Note that the thresholds are not necessarily set independent of income: they can be defined based on the income level of the population group. This way the concern for the poor versus rich could be incorporated. If the coverage rates will also differ with respect to quintiles then  $p_{iq}$  would denote the coverage rate for treatment  $i$  for quintile  $q$ .  $y_{iq}$  is a binary variable that takes a value of 0 if treatment  $i$  leads to catastrophic financial burden in quintile  $q$  and 1 otherwise. Similarly,  $a_{ikq}$  is a binary variable that becomes 1 if  $p_{iq} = P_k$  and 0, otherwise. The resulting model is as follows. We call it the quintile model (Q).

Model Q

$$\begin{aligned} &Max \sum_{i \in I} \sum_{q \in Q} b_i n_{iq} p_{iq} \\ &Max \sum_{i \in I} \sum_{q \in Q} n_{iq} y_{iq} \\ &s.t. \sum_{i \in I} \sum_{q \in Q} c_i n_{iq} p_{iq} \leq B \end{aligned} \tag{17}$$

$$p_{iq} = \sum_{k \in M} a_{ikq} P_k \quad \forall i \in I, \forall q \in Q \tag{18}$$

$$\sum_{k \in M} a_{ikq} \leq 1 \quad \forall i \in I, \forall q \in Q \tag{19}$$

$$(1 - p_{iq})c_i - T_q \leq M(1 - y_{iq}) \quad \forall i \in I, \forall q \in Q \tag{20}$$

$$a_{ikq} \text{ binary } \forall i \in I, \forall k \in M, \forall q \in Q \tag{21}$$

$$y_{iq} \text{ binary } \forall i \in I, \forall q \in Q \tag{22}$$

## Appendix B Malawi data

This data is based on [Ochalek et al., 2018], who proposed an analytical framework to guide the design of health benefit packages. The authors illustrated their framework on Malawian data, which includes estimates of costs and effects of 67 treatments. These estimates are from the Tufts

Global Health Cost-Effectiveness Registry and World Health Organization CHOosing Interventions that are Cost-Effective (WHO-CHOICE) analyses, as well as unpublished data on drug and supply costs and the size of the eligible patient population. Excluding treatments that have very small number of cases or negative net Disability adjusted life years (DALYs) averted, we obtain a set of 46 treatments, whose data is shown in the following table.  $b_i$  is the net DALYs averted per person for treatment  $i$ ,  $n_i$  is number of cases (in 1000s) for treatment  $i$  and  $c_i$  is the cost per person (in \$s) for treatment  $i$ .

<b>Treatment</b>	<i>i</i>	<i>n<sub>i</sub></i>	<i>b<sub>i</sub></i>	<i>c<sub>i</sub></i>
Male Circumcision	1	4073	6.24	36.03
Mgt. Obstructed Labour	2	92	22.02	11.96
Isoniazid for HIV+TB	3	55	25.42	1.45
First line trt. for new TB (adults)	4	14	71.57	12.71
First line trt. for new TB (children)	5	12	70.92	9.75
Mgt of pre-eclampsia	6	20	24.15	2.25
Clean practices and immediate essential NB care	7	671	0.34	0.62
Households owning at least one ITN/LLIN	8	6752	0.03	2.03
Caeserian section	9	34	4.62	19.76
Mass media	10	16879	0.01	0.45
Labour and delivert mgt.	11	918	0.15	1.4
PMTCT of HIV	12	53	2.45	11.32
First line trt. for retreatment TB (adults)	13	2	62.50	50
Ceaserian section (with complication)	14	5	22.20	34.4
First line trt. for retreatment TB (children)	15	2	53.00	33
Malaria:first trimester-uncomplicated	16	305	0.33	3.36
Malaria:second trimester-uncomplicated	17	305	0.33	0.77
Voluntary counselling and testing	18	8031	0.01	4.52
Tetanus toxoid (pregnant women)	19	918	0.10	0.13
Measles vaccine	20	651	0.14	0.81
Rotavirus vaccine	21	651	0.12	4.76
Antenatal care (four visits)	22	918	0.07	12.23
Malaria trt: uncomplicated (adult,<36 kg)	23	4372	0.01	0.79
Malaria trt: uncomplicated (adult,>36 kg)	24	4372	0.01	0.98
Malaria trt: uncomplicated-second line (adult, >36 kg)	25	4372	0.01	0.27
Malaria trt: uncomplicated-second line (adult, <36 kg)	26	4372	0.01	0.14
Vaginal delivery, skilled attendance	27	918	0.06	5.64
Isoniazid for children in contact with pats with TB	28	2	22.00	3.5
Interventions focused on men who have sex with men	29	34	1.18	36.94
Pregnant women sleeping under an ITN	30	1469	0.03	2.04
Newborn sepsis	31	81	0.46	5.15
Mgt of severe malnutrition	32	51	0.71	47.78
Vitamin A supplementation in pregnant women	33	124	0.24	1.01
Antenatal corticosteroids	34	165	0.17	2.46
Interventions focused on female sex workers	35	23	1.22	28.48
Malaria trt: uncomplicated (children,<15kg)	36	1042	0.01	4.39
Malaria trt: uncomplicated (children,>15kg)	37	1042	0.01	4.58
Malaria trt: uncomplicated-second line (children,<15kg)	38	1042	0.01	0.03
Malaria trt: uncomplicated-second line (children,>15kg)	39	1042	0.01	0.07
Under five children who slept under ITN/LLIN	40	494	0.03	2.04
Schistosomiasis mass drug adm.	41	389	0.03	0.2
Antibiotics for pPRoM	42	64	0.16	0.61
Blood safety	43	40	0.23	40.65
Vaginal delivery, with complication	44	138	0.06	5.83
Maternal sepsis case mgt.	45	64	0.11	42.67
Malaria trt: pregnant women-complicated	46	16	0.31	8.75

## Appendix C Algorithm

The epsilon constraint method to obtain solutions to our biobjective programming problems is as follows:

Let the generic model be:

$$\begin{aligned} & \max \text{ Health} \quad \max \text{ Protection} \\ & x \in \mathcal{X} \end{aligned}$$

where  $x$  and  $\mathcal{X}$  denote the decision variable vector and the feasible region of our models, respectively. *Health* is the function used for measuring total health benefit (in our models this is  $\sum_i b_i n_i p_i$  and *Protection* is the function used for measuring financial protection (or exposure) (in our first models this is  $\sum_i n_i y_i$  and in the second one it is  $-\sum_i n_i f_i$ ).

---

### Algorithm 1 Epsilon Constraint Method (*stepsize*)

---

**Initialize:** Set  $\epsilon = 0$ ,  $feasible = 1$ ,  $S = \emptyset$ .  $S$  is the set of solutions.

**while**  $feasible=1$  **do**

Solve P1:

$$\begin{aligned} & \max \text{ Protection} \\ & s.t \text{ Health} \geq \epsilon \\ & x \in \mathcal{X} \end{aligned}$$

**if** The problem is infeasible **then**

$feasible = 0$

**else**

Let the health and financial protection values of the optimal solution be  $H^*$  and  $FP^*$ , respectively. Solve P2:

$$\begin{aligned} & \max \text{ Health} \\ & s.t \text{ Protection} \geq FP^* \\ & x \in \mathcal{X} \end{aligned}$$

Let the health and financial protection values of the optimal solution be  $H^{**}$  and  $FP^{**}$ , respectively.

$S = S \cup (H^{**}, FP^{**})$ .

Set  $\epsilon = H^{**} + stepsize$

**end if**

**end while**

---

*stepsize* is a predetermined parameter that controls movement along the Pareto frontier. The algorithm starts with finding the solution that has the maximum financial protection by solving

P1, then a second model (P2) is solved to find the solution that has maximum health benefit level among the ones with the same financial protection value. If there are alternative solutions that result in the same financial protection score, solving P2 will return the one with the maximum health benefit. The solution of P2 is a Pareto solution and added to the solution set  $S$ . In the next iteration, the algorithm finds the next Pareto solution with a larger *Health* level by adjusting the parameter  $\epsilon$  used in the constraint in P1.

At each iteration, the algorithm finds a Pareto solution by solving P1 and P2 consecutively and terminates when the other end of the Pareto frontier (the Pareto solution with the best health benefit level and the worst financial protection (exposure) level) is reached. After this solution is found and  $\epsilon$  is adjusted accordingly, the modified model P1 will be infeasible and the while loop will terminate.

## Appendix D Discussion on the shape of the Pareto frontier

The shape of the Pareto frontier is data dependent. The projects with core values of 1 (must do projects) and 0 (do not do projects) may define the shape of the Pareto frontier. If the number of such treatments is high, and the projects are similar in terms of size (number of beneficiaries, cost, level of health benefit) one would have a smooth (evenly spaced) Pareto as all nondominated portfolios will have the same decision on such projects (cover fully or do not cover at all).

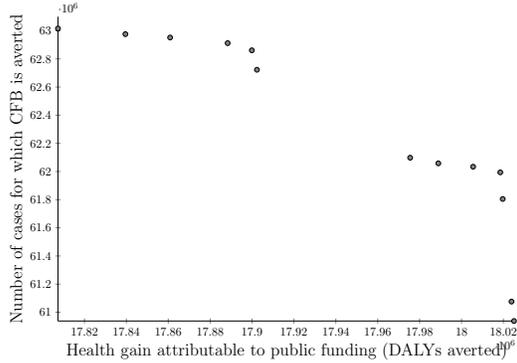
If, however, there are a small number of “large” projects and several “small” projects, which are included in some solutions while being left out at others, there may be sudden jumps at the frontier. One such example occurs at the (90M,5) Pareto (see Figure 10(a) below, which is derived from Figure 6). The third solution (from the maximum DALYs averted edge), is located above the second one and indicates a relatively large improvement in financial protection with very low sacrifice from the total health. This is due to a relatively large treatment (treatment 27 with 918 cases) with a cost-per-person value above the threshold (\$5.64), which is not included in the second solution. When this treatment is subsidized (increasing coverage rate from 0 to 0.25), the number of people protected increases. The additional cost is compensated by reducing coverage rates of

already easier-to-pay treatments such as treatment 30 (which has a per person cost of only \$2.0 ). The second jump occurs due to another large treatment (22 with again 918 patients, with relatively large per-person cost of \$12.23) being included in the package (in solution 8). Note also that these jumps occur for all Pareto frontiers summarized in Figure 6 but in the ones corresponding low budgets (70M,3) and (70M,5) they are associated with more sacrifice from health gain per unit of financial protection (this is a direct result of the budget being scarce).

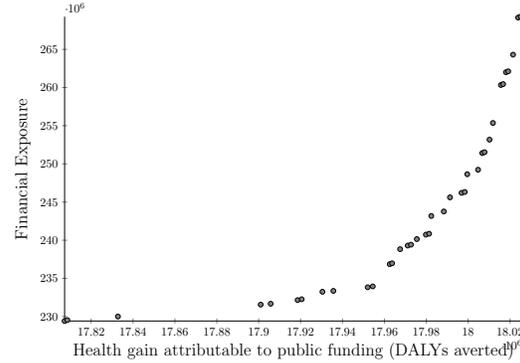
The appearance of these jumps in model B solutions also has something to do with the fact that in the first model (Model B), we “count” the number of cases for which the expenses are below the threshold. Any financial burden is simply categorized as catastrophic or not, there is no measure of how catastrophic. The method does not distinguish two solutions with the same number of cases averted, with respect to how much total catastrophic expenditure is born.

As mentioned in section 3, the second approach (Model C) allows us to go beyond simply designating a financial burden as being catastrophic or not and to incorporate different assumptions about “how catastrophic”. It attempts to measure it and finds a smoother Pareto frontier with more intermediate solutions, in which treatments leading to very high out-of-pocket expenditures are also funded to various degrees (these intermediate coverage rates for expensive treatments may not be seen in Model B solutions, as the out-of-pocket payments will already stay at the higher side of the threshold). For example, treatment 43 has a per-person cost of \$40.65. In Model B solutions, it either has a coverage rate of 0.9 (which means the cases are saved from catastrophic burden, given a threshold value of 5) or 0. In Model C, there are Pareto solutions with coverage rates of 0.9, 0.4, 0.25 and 0 for that treatment. In that sense, the second model explores a larger spectrum of options.

This is also mirrored by the Pareto frontiers, as seen in Figures 10(a) and 10(b). Model C has a smoother Pareto with more solutions.



(a) Pareto solutions of Model B (90M, 5 coverage rates). Note that both criteria are of a maximization type.



(b) Pareto solutions of Model C (90M, 5 coverage rates). Note that financial exposure is to be minimized, hence the shape of the Pareto curve is different.

Figure 10: Pareto frontiers of Models B and C for (90M, 5 coverage rates)

## References

[Chalkidou et al., 2016] Chalkidou, K., Glassman, A., Marten, R., Vega, J., Teerawattananon, Y., Tritasavit, N., Gyansa-Lutterodt, M., Seiter, A., Kieny, M. P., Hofman, K., and Culyer, A. J. (2016). Priority-setting for achieving universal health coverage. *Bulletin of the World Health Organization*, 94(6):462–467.

[Crown et al., 2018] Crown, W., Buyukkaramikli, N., Sir, M. Y., Thokala, P., Morton, A., Marshall, D. A., Tosh, J. C., Ijzerman, M. J., Padula, W. V., and Pasupathy, K. S. (2018). Application of constrained optimization methods in health services research: Report 2 of the ispor optimization methods emerging good practices task force. *Value in Health*, 21(9):1019–1028.

[Crown et al., 2017] Crown, W., Buyukkaramikli, N., Thokala, P., Morton, A., Sir, M. Y., Marshall, D. A., Tosh, J., Padula, W. V., Ijzerman, M. J., Wong, P. K., and Pasupathy, K. S. (2017). Constrained optimization methods in health services research—an introduction: Report 1 of the ispor optimization methods emerging good practices task force. *Value in health : the journal of the International Society for Pharmacoeconomics and Outcomes Research*, 20(3):310–319.

[Ehrgott, 2005] Ehrgott, M. (2005). *Multicriteria optimisation*. Springer, Berlin.

- [Glassman et al., 2017] Glassman, A., Giedion, U., and Smith, P. C. (2017). *What's in, what's out: designing benefits for Universal Health Coverage*. Center for Global Development, Washington DC.
- [Hauck et al., 2019] Hauck, K., Morton, A., Chalkidou, K., Chi, Y. L., Culyer, A., Levin, C., Meacock, R., Over, M., Thomas, R., Vassall, A., Verguet, S., and Smith, P. C. (2019). How can we evaluate the cost-effectiveness of health system strengthening? a typology and illustrations. *Social Science & Medicine*, 220:141–149.
- [Jamison and DCP3, 2018] Jamison, D. T. and DCP3 (2018). Universal health coverage and intersectoral action for health: key messages from disease control priorities, 3rd edition. *Lancet*, 391:1108–1120.
- [Karsu, 2013] Karsu, O. (2013). *Using holistic multicriteria assessments: the convex cones approach*. Wiley, New York.
- [Korhonen et al., 1984] Korhonen, P., Wallenius, J., and Zionts, S. (1984). Solving the discrete multiple criteria problem using convex cones. *Management Science*, 30(11):1336–1345.
- [Megiddo et al., 2016] Megiddo, I., Colson, A., Chisholm, D., Dua, T., Nandi, A., and Laxminarayan, R. (2016). Health and economic benefits of public financing of epilepsy treatment in india: An agent-based simulation model. *Epilepsia*, 57(3):464–74.
- [Megiddo et al., 2014] Megiddo, I., Colson, A. R., Nandi, A., Chatterjee, S., Prinja, S., Khera, A., and Laxminarayan, R. (2014). Analysis of the universal immunization programme and introduction of a rotavirus vaccine in india with indiasim. *Vaccine*, 32 Suppl 1:A151–61.
- [Megiddo et al., 2018] Megiddo, I., Klein, E., and Laxminarayan, R. (2018). Potential impact of introducing the pneumococcal conjugate vaccine into national immunisation programmes: an economic-epidemiological analysis using data from india. *BMJ global health*, 3(3):e000636.
- [Morton, 2014] Morton, A. (2014). Aversion to health inequalities in healthcare prioritisation: A multicriteria optimisation perspective. *Journal of Health Economics*, 36:164–173.

- [Morton et al., 2018] Morton, A., Arulselvan, A., and Thomas, R. (2018). Allocation rules for global donors. *Journal of Health Economics*, 58:67–75.
- [Morton et al., 2016a] Morton, A., Keisler, J., and Salo, A. (2016a). *Multicriteria Portfolio Decision Analysis for Project Selection*. Springer, New York.
- [Morton and Lauer, 2017] Morton, A. and Lauer, J. A. (2017). *Comparing apples and oranges: strategies to weigh health against other social values*. Center for Global Development, Washington DC.
- [Morton et al., 2016b] Morton, A., Thomas, R., and Smith, P. C. (2016b). Decision rules for allocation of finances to health systems strengthening. *Journal of Health Economics*, 49:97–108.
- [Nandi et al., 2016] Nandi, A., Colson, A. R., Verma, A., Megiddo, I., Ashok, A., and Laxminarayan, R. (2016). Health and economic benefits of scaling up a home-based neonatal care package in rural india: a modelling analysis. *Health policy and planning*, 31(5):634–44.
- [Ochalek et al., 2018] Ochalek, J., Revill, P., Manthalu, G., McGuire, F., Nkhoma, D., Rollinger, A., Sculpher, M., and Claxton, K. (2018). Supporting the development of a health benefits package in malawi. *BMJ global health*, 3(2):1–14.
- [Salo et al., 2011] Salo, A., Keisler, J., and Morton, A. (2011). *Portfolio Decision Analysis: methods for improved resource allocation*. Springer, New York.
- [Smith, 2013] Smith, P. C. (2013). Incorporating financial protection into decision rules for publicly financed healthcare treatments. *Health Economics*, 22:180–193.
- [Smith and Yip, 2016] Smith, P. C. and Yip, W. N. (2016). The economics of health system design. *Oxford Review of Economic Policy*, 32(1):21–40.
- [van Baal et al., 2016] van Baal, P., Meltzer, D., and Brouwer, W. (2016). Future costs, fixed health care budgets and the decision rules of cost effectiveness analysis. *Health Economics*, 25(2):237–248.

- [van Baal et al., 2018] van Baal, P., Morton, A., and Severns, J. L. (2018). Health care input constraints and cost effectiveness analysis decision rules. *Social Science & Medicine*.
- [Verguet et al., 2016] Verguet, S., Kim, J. J., and Jamison, D. T. (2016). Extended cost-effectiveness analysis for health policy assessment: A tutorial. *PharmacoEconomics*, 34(9):913–23.
- [Verguet et al., 2015] Verguet, S., Laxminarayan, R., and Jamison, D. T. (2015). Universal public finance of tuberculosis treatment in india: an extended cost-effectiveness analysis. *Health economics*, 24(3):318–32.
- [Verguet et al., 2013] Verguet, S., Murphy, S., Anderson, B., Johansson, K. A., Glass, R., and Rheingans, R. (2013). Public finance of rotavirus vaccination in india and ethiopia: an extended cost-effectiveness analysis. *Vaccine*, 31(42):4902–10.
- [WHO, 2010] WHO (2010). *World Health Report 2010: Financing for Universal Health Coverage*. World Health Organization, Geneva.
- [Winston, 2004] Winston, W. L. (2004). *Operations Research: Applications and Algorithms*. Thomson/ Brooks/Cole, 4 edition.