

Title page

Title: Applying the “no-one worse off” criterion to design Pareto efficient HIV responses in Sudan and Togo

Authors: Robyn M. Stuart^{1,2,§}, Hassan Haghparast-Bidgoli³, Jasmina Panovska-Griffiths^{4,5}, Laura Grobicki³, Jolene Skordis³, Cliff C. Kerr^{1,6}, David J. Kedziora^{1,3,6,7}, Rowan Martin-Hughes¹, Sherrie L. Kelly^{1,7}, David P. Wilson^{1,7}

Author affiliations:

- 1 Burnet Institute, Melbourne, Australia
 - 2 Department of Mathematical Sciences, University of Copenhagen, Copenhagen, Denmark
 - 3 Institute for Global Health, University College London, London, UK
 - 4 Departments of Applied Health Research and Mathematics, University College London, London, UK
 - 5 Department of Global Health and Development, London School of Hygiene and Tropical Medicine, London, UK
 - 6 School of Physics, University of Sydney, Sydney NSW, Australia
 - 7 Monash University, Melbourne, Australia
- § Corresponding author: Robyn M. Stuart

Author e-mail addresses:

RMS: robyn@math.ku.dk
HH-B: h.haghparast-bidgoli@ucl.ac.uk
JP-G: j.panovska-griffiths@ucl.ac.uk
LG: laura.grobicki@ucl.ac.uk
JS-W: j.skordis@ucl.ac.uk
CCK: cliff@optimamodel.com
DJK: david.kedziora@burnet.edu.au
RM-H: rowan.martin-hughes@burnet.edu.au
SLK: sherrie.kelly@burnet.edu.au
DPW: david.wilson@burnet.edu.au

Abstract

Introduction: Globally, there is increased focus on getting the greatest impact from available health funding. However, the pursuit of overall welfare maximization may mean some are left worse off than before. Pareto efficiency takes welfare shifts into account by ruling out funding reallocations that worsen outcomes for any person or group.

Methods: Using the Optima HIV model, studies of HIV response efficiency were conducted in Sudan in 2014 and Togo in 2015. In this paper we estimate the welfare maximizing and Pareto efficient allocations for these two national HIV budgets, using data from the original studies.

Results: We estimate that if the 2013 HIV budget for Sudan was annually available to 2020 but with funds reallocated according to the welfare maximizing allocation, a 36% reduction in cumulative new infections could be achieved between 2014–2020. We also find that this is Pareto efficient. In Togo, however, we find that it is possible to reduce overall new infections but applying the Pareto efficiency criterion means that shifts in emphases cannot occur in the HIV response without additional resources.

Discussion: Protecting service coverage for key population groups is not necessarily equivalent to protecting health outcomes. In some cases, requiring Pareto efficiency may reduce the potential for population-wide welfare gains, but this is not always the case.

Conclusion: Pareto efficiency may be an appropriate addition to the quantitative toolset for evaluating HIV responses.

Keywords: key populations, HIV prevention, HIV modeling, allocative efficiency, Pareto efficiency, Pareto optimality

Introduction

The efficient allocation of health resources is an important objective for health systems as populations age, the global burden of non-communicable disease grows, and health financing stagnates. Broadly speaking, efficiency measures whether the best outcomes are being obtained for a given budget[1], but to make use of this very broad definition it is necessary to understand what is meant by “best outcomes”, and what barriers might exist to obtaining them. Efficiency studies often focus on quantitatively answering the latter question[2], but the former is more subjective, and even if we agree on a way to quantify an individual’s health outcomes, a separate set of complications arises when it comes to aggregating across individuals to determine net social welfare[3, 4].

Quantifying social welfare becomes pertinent when considering *allocative efficiency*, i.e. evaluating how resources are allocated across interventions. A healthcare system can generally be considered allocatively efficient if no other allocation would increase overall population welfare. However, the pursuit of a welfare maximizing outcome may mean that some are left worse off than they were before. *Pareto efficiency* or *Pareto optimality* is an alternative efficiency criterion that considers these welfare shifts: an allocation is *Pareto efficient/optimal* if there is no other allocation that would make one individual better off without making at least one other individual worse off. (A third efficiency criterion, Kaldor-Hicks efficiency, combines aspects of both these criteria by allowing for hypothetical compensation of those made worse off. Thus, a welfare-maximizing outcome could be Kaldor-Hicks efficient even if not Pareto efficient. However, Kaldor-Hicks efficiency requires the assumption that utility is transferable, which is not the case for health.) Pareto

efficiency is rarely used in health economics, because generally speaking, any reallocation of health resources will result in at least one person being worse off[5]. We therefore propose to consider a response Pareto efficient if no defined demographic or risk group can be made better off without some other group being made worse off. This population-group definition is less restrictive than an individual-level definition and potentially more useful, given that evaluation of health responses often considers the impact on certain population cohorts, particularly those at greater risk.

In this study, we compare two allocation strategies – welfare maximizing and Pareto efficiency – for evaluating HIV responses. We propose that it may be appropriate to aim for ‘group Pareto efficiency’ in the context of infectious diseases. We begin by presenting some theory for Pareto efficiency for infectious diseases. Next, we investigate HIV programmatic priority setting using two country examples, based on HIV allocative efficiency studies conducted in Sudan[6] and Togo[7]. To investigate allocation strategies, we use Optima HIV, a mathematical model designed to facilitate estimation of the welfare maximizing allocation of HIV resources. We describe how the model’s optimization algorithm was extended to estimate Pareto efficient allocations. We show that in Sudan, the welfare maximizing allocation was Pareto efficient, but in Togo the welfare-maximizing allocation was not Pareto efficient.

Methods

Theory of Pareto efficiency in infectious diseases

To outline the key theoretical aspects, we consider a simplified epidemic context with two programs, one targeted at female sex workers (FSW) and the other at people who inject drugs (PWID). A total budget of \$2m is split equally between the two programs. The number of new HIV infections acquired by FSW and PWID depends how much of the budget is

allocated to each program. Figure 1 presents two possible scenarios. In the first (Figure 1a), there is limited mixing between FSW and PWID, so reducing new infections in one population does not greatly affect the other. In the second (Figure 1b), injection is the primary mode of HIV transmission, and infections acquired by PWID are sexually transmitted to FSW. In this instance, it benefits both populations to shift more funding towards the PWID program.

[Figure 1 goes here]

Scenarios like those presented in Figure 1b distinguish the theory of Pareto efficiency in infectious diseases from how it is usually applied in welfare economics, where it is assumed that people's utility increases monotonically with the share of resources they are allocated. This implies that Pareto improvements are only possible via technical efficiency improvements, increasing budgets, or utility transfers. However, for communicable diseases, preventing a primary case in one population may avert multiple downstream infections due to the dynamic nature of pathogen transmission. Pareto efficiency has been considered in theoretical infectious disease models[8-10], but not extensively in real-world settings.

Using Optima HIV to determine welfare-maximizing and Pareto efficient allocations

Optima HIV is an epidemic model linked to a programmatic response module that can estimate the optimal allocation of a total HIV budget across different programs for a given objective and set of intervention cost functions[11]. The entire population can be divided into as many subpopulations as required to capture variation in vulnerability to HIV infection. We consider an example where the population P is divided into N subpopulations, denoted P_i ($i=1, \dots, N$), and the HIV response has M programs with annual budgets B_j ($j=1, \dots, M$). We call $[B_1, \dots, B_M]$ allocation \mathbf{A} . The health outcomes of population P_i depend on \mathbf{A} , so we write

$O_i(A)$ where O stands for outcomes. The health outcome of the entire population is the sum of outcomes in each subpopulation, i.e., $O(\mathbf{A}) = \sum_{i=1}^N O_i(\mathbf{A})$.

To determine the welfare-maximizing allocation, we choose the health outcome O of interest (typically new HIV infections, DALYs, deaths, or another country-specific outcome), and find the allocation \mathbf{A}^* that solves the optimization problem $\min_{\mathbf{A}} O(\mathbf{A})$, subject to the constraint that the overall budget $B = \sum_{j=1}^M B_j$ remains constant.

Extending this to include Pareto efficiency requires the additional constraint:

$$\min_{\mathbf{A}} O(\mathbf{A}) \text{ subject to } O_i(\mathbf{A}) \leq O_i(\mathbf{A}_B),$$

where \mathbf{A}_B is the baseline allocation. It is easy to show that the \mathbf{A}^* that solves this is Pareto efficient, since the constraint combined with the minimization ensures that it is welfare-maximizing with no subpopulation worse off.

Overview of HIV allocative efficiency studies conducted in Sudan and Togo

In 2014, Sudan's national AIDS program reviewed its National Strategic Plan (NSP) on HIV and prepared a concept note for the Global Fund for AIDS, Tuberculosis and Malaria (the Global Fund), and in 2015, the Government of Togo underwent the same process. Each country approached the World Bank with a request to conduct allocative efficiency analyses to inform the prioritization of their HIV responses and the development of their Global Fund concept notes. In both cases, analyses focused on how HIV resources should be allocated for greatest epidemiological impact. The Optima HIV model was applied in both country's processes .

Evaluating funding allocations in Sudan and Togo

In Sudan, the primary objective of the study was to minimize new infections over 2014-2020, which meant determining the welfare-maximizing allocation of the HIV budget. In Togo, the primary objective was to minimize a combination of deaths and new infections over 2015-2020 whilst ensuring that infections did not increase in any subpopulation. This meant

exploring whether there were allocations that were Pareto superior to the 2014 allocation. For this study, we estimated both the welfare-maximizing and Pareto efficient allocation for both countries.

Results

In Sudan, we found that even if overall investment in HIV does not increase, there is a welfare maximizing allocation that could attain a 36% reduction in cumulative new infections over 2014-2020 (Figure 2b). This reallocation is also estimated to reduce infections in each subpopulation and is therefore Pareto efficient.

In contrast, in Togo we found that the latest reported allocation of investments was already close to welfare maximizing (Figure 2b). However, the welfare-maximizing allocation implied a 7% increase in cumulative new infections among males aged 15-24 years (despite a 5% reduction in new infections overall), and thus is not Pareto efficient. We found no Pareto-improving allocations of the HIV budget in Togo, indicating that any shift from status quo investment patterns would make at least one subpopulation worse off.

Complex epidemic dynamics mean that shifting funds away from a particular population may not translate to an increase in infections in that population. This can be observed for Sudan, where new infections among the general population were projected to decrease despite defunding prevention programs targeted at this population (effectively a kind of compensation: those who gain from reallocation 'compensate' others by lowering their HIV acquisition risk). Likewise, mother-to-child transmission (MTCT) was projected to decline even if funding for prevention of MTCT were to decrease. In both cases, this is due to the

large proportion of new infections (~1/3 in 2014) contributed by sex work[6], and the relative cost-effectiveness of preventing infections among sex workers.

Discussion

In the context of infectious diseases, preventing primary cases may avert multiple downstream infections, so programs that are most effective at averting primary cases are not necessarily those that will maximize overall population health. Mathematical models have proven useful for estimating the cumulative downstream effects of responses[12]. In this study, we proposed an extension for how models can be used, namely as a tool for estimating the scope for Pareto improvements.

One limitation of this study is that the type of Pareto efficiency we have considered is not the typical individual-level definition, and is thus sensitive to the subpopulations considered. For example, the welfare-maximizing allocation in Togo was characterized as Pareto inferior to the baseline allocation because it meant an increase in infections among males aged 15–24 years. However, had a different population grouping been defined for analysis (e.g. grouping all males aged 15-49), the welfare-maximizing allocation would have been Pareto efficient. If the subpopulations are small enough, it may be impossible to find a Pareto superior allocation. We also do not consider potential redistributions of health within subpopulations. However, within HIV there are standard key populations (e.g. FSW, MSM, and PWID) of interest, and protecting the health outcomes of these groups could be done fairly consistently. The existence of a Pareto superior allocation also depends on the programs considered: if programs are divided into intervention implementation modalities, finding a Pareto optimal allocation may be easier. We also note the limitations of Pareto efficiency itself as a social choice rule. Its drawbacks (including the disregard of equality, social welfare, or resource wastage) are well-known, and assuming that status quo health outcomes are worth protecting

may not always be reasonable. Other factors, such as equity and fairness, need to be considered to overcome these limitations.

Allocative efficiency studies have been conducted in numerous settings over the past decade by our group and others [13-17]. Whilst decision-makers have recognized the importance of optimally allocating funding to get the best possible health outcomes, it is still common to incorporate ethical, logistical, and political constraints to protect the interests of various groups and ensure that modeling recommendations are practicable. However, protecting intermediary health measures (such as program coverage) may not be the most effective way to protect the health outcomes. Thus, Pareto efficiency may be a useful addition to the standard set of tools for evaluating responses to infectious diseases, even though it cannot be the sole criterion by which distributional fairness can be assessed.

Acknowledgements

Author contributions:

RMS: conceived study, wrote Pareto efficiency algorithm, conducted Sudan analyses and wrote first draft
HH-B: conducted Togo analyses, contributed to discussion on social optimality
JP-G: conducted Togo analyses
LG: conducted Togo analyses
JS-W: supervised Togo analyses
CCK: reviewed Pareto efficiency algorithm
DJK: technical assistance to Togo and code development
RM-H: updated Sudan and Togo analyses
SLK: updated Sudan and Togo analyses
DPW: conceptual guidance/supervision
All: contributed to paper editing

Funding: World Bank and Australian NHMRC. JPG's research was supported by the National Institute for Health Research (NIHR) Collaboration for Leadership in Applied Health Research and Care North Thames at Bart's Health NHS Trust (NIHR CLAHRC North Thames). The views expressed in this article are those of the author(s) and not necessarily those of the NHS, the NIHR, or the Department of Health and Social Care in the UK.

References

1. Williams, A., *Priority setting in public and private health care. A guide through the ideological jungle*. J Health Econ, 1988. 7(2): p. 173-83.
2. Chisholm, D. and D. Evans, *Improving health system efficiency as a means of moving towards universal coverage. Background Paper No 28.*, in *World Health Report*, W.H. Organisation, Editor. 2010.
3. Theil, H., *Linear aggregation of economic relations*. 1954: North-Holland Pub. Co.
4. Arrow, K.J., *A Difficulty in the Concept of Social Welfare*. Journal of Political Economy, 1950. 58: p. 328-346.
5. Palmer, S. and D.J. Torgerson, *Definitions of efficiency*. BMJ, 1999. 318(7191): p. 1136.
6. The World Bank, *Sudan's HIV response: value for money in a low-level HIV epidemic*. 2014.
7. The World Bank, *Optimizing Investments in Togo's HIV Response*. 2016.
8. Klepac, P., R. Laxminarayan, and B.T. Grenfell, *Synthesizing epidemiological and economic optima for control of immunizing infections*. Proceedings of the National Academy of Sciences of the United States of America, 2011. 108(34): p. 14366-14370.
9. Philipson, T.J., *Economic epidemiology and infectious diseases*, in *Handbook of Health Economics*, A.J.C.J.P. Newhouse, Editor. 2000, Elsevier.
10. Lundström, N.L.P., H. Zhang, and Å. Brännström, *Pareto-efficient biological pest control enable high efficacy at small costs*. 2018: <https://arxiv.org/pdf/1608.00603.pdf>.
11. Kerr C.C., et al., *Optima: A Model for HIV Epidemic Analysis, Program Prioritization, and Resource Optimization*. J Acquir Immune Defic Syndr, 2015. 69(3): p. 365-76.
12. Eaton, J.W., et al., *Health benefits, costs, and cost-effectiveness of earlier eligibility for adult antiretroviral therapy and expanded treatment coverage: a combined analysis of 12 mathematical models*. Lancet Glob Health, 2014. 2(1): p. e23-34.
13. Stuart, R.M., et al., *How should HIV resources be allocated? Lessons learnt from applying Optima HIV in 25 countries*. 2017.
14. The World Bank, *Scaling up HIV treatment for MSM in Bangkok: what does it take? – a modelling and costing study*. 2015.
15. Kahn, J., et al., *Using Models to Guide HIV/AIDS Policy: A Synthesis of Current Models to Determine Resource Allocation Cost-Effectiveness*, in *Disease Control Priorities*, K. Holmes, et al., Editors. 2016, World Bank: Washington, DC. p. 39.
16. Anderson, S.-J., et al., *Maximising the effect of combination HIV prevention through prioritisation of the people and places in greatest need: a modelling study*. Lancet, 2014. 384(9939): p. 249-256.
17. McGillen, J.B., et al., *Optimum resource allocation to reduce HIV incidence across sub-Saharan Africa: a mathematical modelling study*. Lancet HIV, 2016. 3(9): p. e441-e448.

Tables and figures

Figure 1. Pareto efficiency and welfare maximization. The green, blue, and black lines represent the number of new HIV infections acquired by female sex workers (FSW), people who inject drugs (PWID), and the total population, respectively. The status quo allocation is for US\$1m to be allocated to the PWID program and US\$1m to the FSW program. Panel (a) illustrates a scenario with limited mixing between the two populations, so reducing new infections in one population doesn't greatly reduce infections in the other, although overall infections can be reduced by shifting funding to PWID programs. Panel (b) illustrates a scenario where shifting funds to PWID also reduces infections in FSW, because the majority of infections in FSW are acquired through sexual contact with PWID.

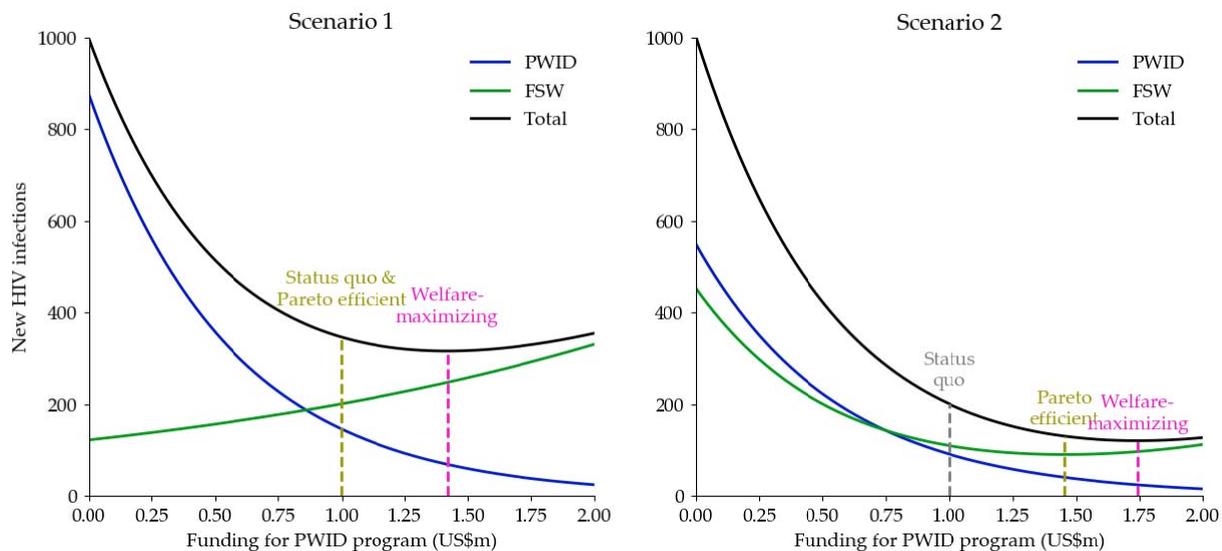
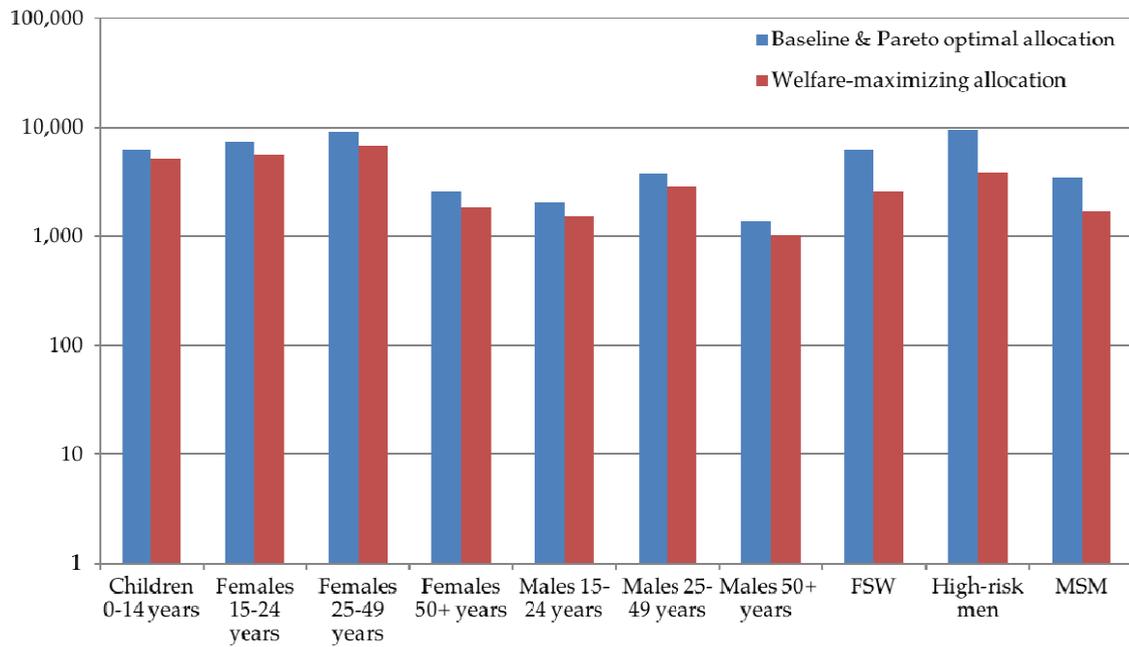


Figure 2. Cumulative HIV infections projected under the baseline, Pareto efficient, and welfare-maximizing allocations of HIV budgets in (a) Sudan and (b) Togo. The allocations are shown in table 1 of the supplementary materials, <http://links.lww.com/QAD/B438>.

A)



B)

