

Doublethink in global prioritisation

Hilary Greaves (Oxford)

22nd AAPAE Conference

University of Auckland, 11 July 2015

Outline

- 1) Global prioritisation: child mortality, family planning and the cancellation worry
- 2) Making it quantitative: the benefit-cost approach
- 3) CBA for child mortality reduction
 - 1) Arguments for not counting 'knock-on effects'
 - 2) Critique of the CBA
- 4) CBA for family planning
 - 1) An excursion into population axiology
 - 2) Critique of the CBA
- 5) Conclusions

Two top priorities? Child mortality and family planning

- The big question of global prioritisation: How best to improve the world?
 - Armed conflict, biodiversity, climate change, water and sanitation, education/illiteracy, malnutrition, infectious disease, natural disasters, corruption, trade barriers, pollution, terrorism, gender equality, financial instability, global economic inequality, environmental sustainability...
- Difficult issues surrounding
 - Which improvements can (reliably) be made
 - How to trade off different types of benefit against one another
- Tough choices: not everything that is *good* is a *top priority*
- Two 'top picks':
 - Reducing child mortality
 - Family planning
 - In tension with one another?

Child mortality

- MDG 4: “Reduce by two thirds, between 1990 and 2015, the under-five mortality rate”
- Global under-5 mortality rates (U5MR):
 - In 1990: 9%
 - In 2013: 4.6%
- This is a partial success story. But there is plenty more progress still to be made...
 - 17,000 children under 5 die every day
 - In 12 countries of sub-Saharan Africa, the U5MR is still over 10%
- Also a ‘top pick’ according to
 - The Copenhagen Consensus
 - GiveWell; Giving What We Can
 - The Bill and Melinda Gates Foundation



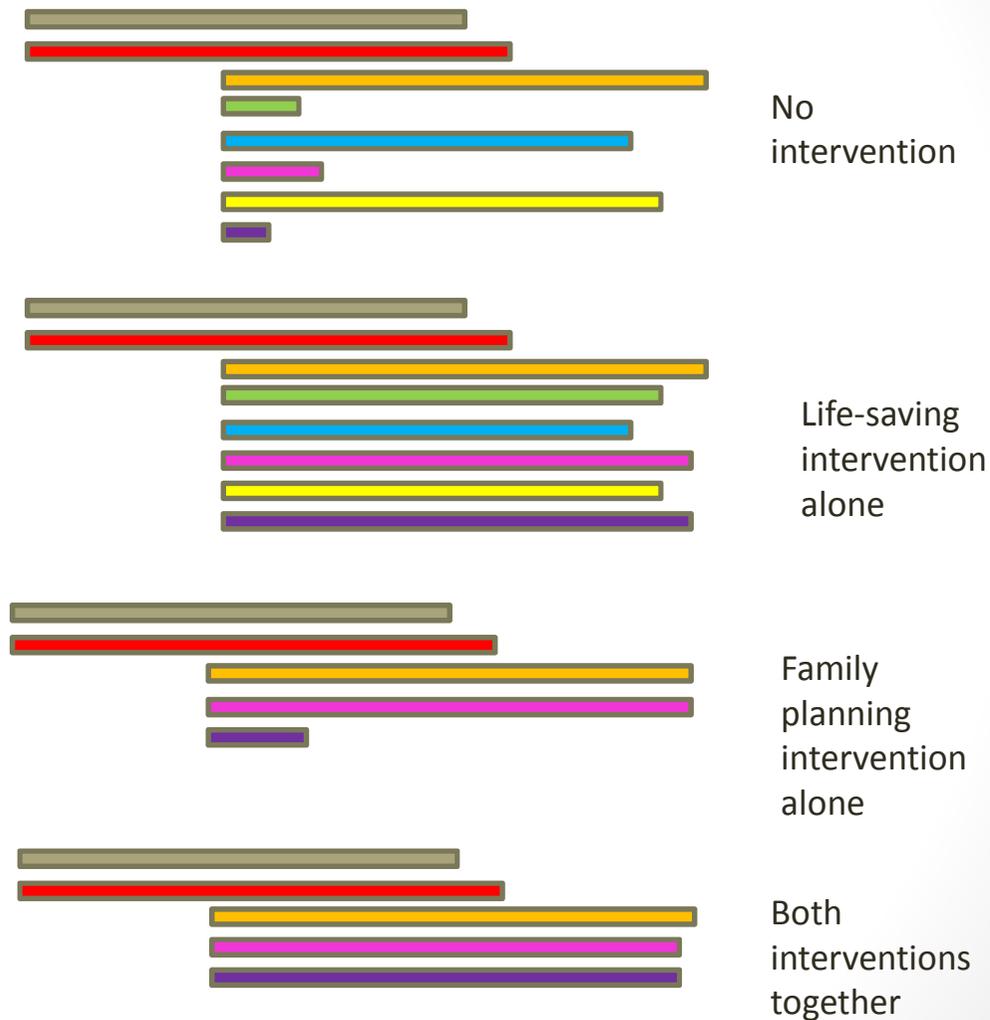
Family Planning

- (Perceived and/or real) benefits of increasing contraceptive use:
 - Feminist concerns:
 - Women's empowerment and gender equality
 - Neo-Malthusian concerns:
 - Environmental sustainability
 - Economic growth
 - More natural resources per capita
 - Health benefits (reduced mortality)
 - Decreased maternal mortality
 - Decreased infant and child mortality (due to increased birth spacing)
- 1960s and 1970s: substantial funding for family planning as a form of international aid (e.g. UNFPA, World Bank)
- More recently, this has become unfashionable... but perhaps should be “back on the agenda”?
 - A ‘top pick’ according to, e.g.
 - The Bill and Melinda Gates Foundation
 - (The Copenhagen Consensus 2012)



The cancellation worry: a first pass

- The net effect of funding *both* infectious disease control *and* family planning, up to permutations of possible people, is to prevent some short lives.
 - *Not* really: to ‘save’ life-years
 - And *not*: to reduce the instantaneous population size
- This may well be good to *some* extent. But *how* good? (Still a ‘top pick’?)



Making it quantitative

- The quantitative question: Granted that e.g. improving sanitation is good, *how* good is it, and how much does it cost?
- Cost-benefit analysis: For each proposed intervention,
 - Estimate the amount of benefit per dollar spent, in monetary terms;
 - Calculate a benefit:cost ratio.
- Interventions targetting child mortality and/or family planning top the prioritisation charts if, and then because, their estimated *benefit:cost ratios* exceed those of other candidate interventions.

PRIORITIZED LIST

	<i>Challenge</i>	<i>Solution</i>
1	Hunger & Education	Bundled Interventions to Reduce Undernutrition in Pre-Schoolers
2	Infectious Disease	Subsidy for Malaria Combination Treatment
3	Infectious Disease	Expanded Childhood Immunization Coverage
4	Infectious Disease	Deworming of Schoolchildren
5	Infectious Disease	Expanding Tuberculosis Treatment
6	Hunger & Biodiversity & Climate Change	R&D to Increase Yield Enhancements
7	Natural Disasters	Investing in Effective Early Warning Systems
8	Infectious Disease	Strengthening Surgical Capacity
9	Chronic Disease	Hepatitis B Immunization
10	Chronic Disease	Acute Heart Attack Low Cost Drugs

Cost-benefit analysis for child mortality reduction

- Valuing reductions in child mortality:
 - Set a monetary value on a healthy life-year. (Usually, for developing countries: about US\$1000.)
 - Measure the benefits of mortality-reduction interventions that combat infectious diseases via the number of healthy life-years that they save.
 - E.g. saving the life of a 5-year-old, in a country with a life expectancy of 65, might be valued at \$60,000.
- Analysis of the work of Against Malaria Foundation (AMF):
 - GiveWell's estimate: One under-5 death is averted per \$3000 donated to AMF.
 - Therefore, if averting an under-5 death saves 60 healthy life-years and each health life-year is valued at \$1,000, then the benefit:cost ratio for donations to AMF is 20:1.
- GiveWell notes explicitly that its calculation *excludes*:
 - Deaths *over age 5* averted;
 - *Non-fatal* cases of malaria averted;
 - *Other mosquito-borne diseases* (other than malaria) prevented.

A tentative critique of the benefit-cost calculations for child mortality reduction

- The calculations just described *do not* take account of ‘indirect’ effects of increasing population (via extending lives) on:
 - Economic development, environmental degradation, political instability, per capita resource shares...
 - Knock-on future increases (or decreases) in the birth rate
- Why not?
 - Negligible in comparison with the direct health effects?
 - Not *obviously* (cf. neo-Malthusianism...).
 - Too hard to estimate?
 - This does not justify pretending that the effects in question are zero.
 - It’s morally inappropriate to take account of indirect effects?...

Reasons not to include indirect effects?

- In each of the following two cases, two patients, A and B, present to the emergency room with life-threatening injuries. Sadly, there are only enough resources to treat one, and the one not treated will die. All possibly-relevant things are equal, *except* that
 - Case 1: Patient A is a surgeon, who, if saved, would go on to save the lives of five others.
 - Case 2: Patient A is the single parent of a young child.
 - Case 3: Patient A is mid-career, while Patient B is retired.
 - Case 4: Patient A is a high earner, while Patient B sweeps the streets.
- Question: Do the positive 'indirect' effects of saving Patient A, in any or all of these cases, make it permissible (or obligatory) for the emergency-room doctor to prioritise A over B?
- Arguments for a negative answer
 - 'Moral' arguments: To prioritise A on the basis of her greater utility to others
 - Would be to 'treat B merely as a means'.
 - Would be 'unfair' to B.
 - Would be to 'fail to respect the equal worth of B as a person'.
 - 'Pragmatic' arguments: Thanks to some contingent feature of the situation, e.g.
 - epistemic limitations of the doctor's situation
 - risk of undermining of doctor-patient trust,any attempt to implement such a prioritisation system would lead to lower welfare overall.

The imperative to consider *all* effects for global prioritisation purposes, however indirect

- Distinguish between three types of decision scenario:
 - Micro-allocation: A doctor, prioritising among patients
 - Mid-level allocation: A health authority, prioritising among healthcare interventions
 - Macro-allocation: A government or philanthropist, prioritising among healthcare *and other* interventions
- A general presumption in favour of counting *all* the effects of a (health) intervention: Indirect and non-health benefits (and harms) are nonetheless real! So it would be irrational not to count them.
- This presumption is particularly strong in the 'macro' context:
 - We don't e.g. restrict to "electricity benefits" when estimating the value of electricity provision; why systematically disadvantage healthcare?
 - Society's *reasons* for supporting a healthcare system are in fact economic as well as more 'intrinsic'.
- And a special exception for health is unmotivated in the 'macro' context:
 - Can't object *here* that "healthcare resources have the aim of promoting health".
 - The pragmatic objections to 'bedside prioritisation' don't apply.
 - Objections based on "treating people as a means" etc. are less likely to apply, since here *we're not directly choosing between individuals* at all.
- For macro purposes (at least), a given 'indirect' effect can be neglected only if, and then because, its expected value is estimated to be negligible. (Brock 2003, Lippert-Rasmussen and Lauridsen 2010)
- 'Doublethink, part I': Ignoring indirect effects *even if they are sizeable*

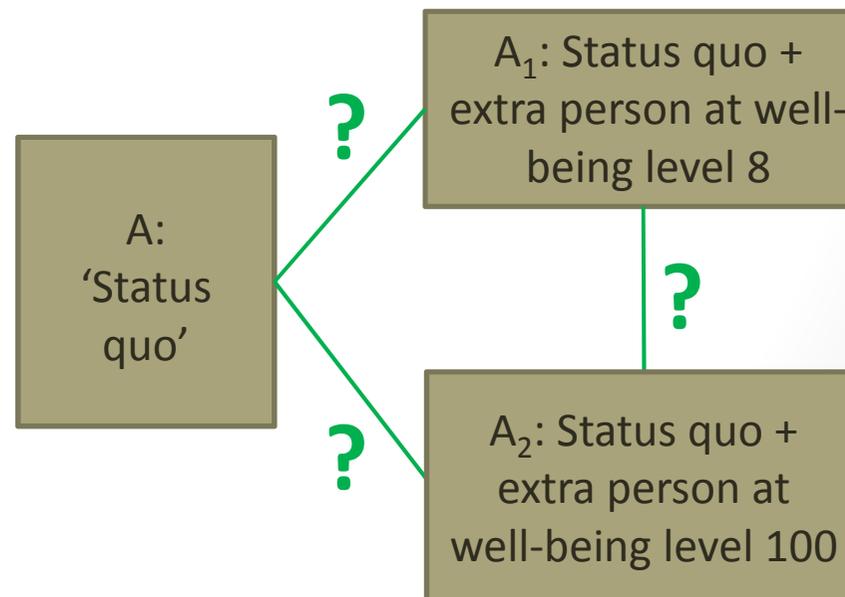
Cost-benefit analysis for family planning

- Kohler's (2012) cost-benefit analysis of family planning interventions:
 - Benefits counted:
 - On averting deaths: Meeting all 'unmet need' would (annually)
 - Cost \$3.6bn
 - Avert
 - 640,000 newborn deaths, saving 65 life-years each
 - 150,000 maternal deaths, saving 37 life-years each
 - @ \$1k per life-year, this is equivalent to \$47bn
 - Benefit:cost ratio of 13:1 from this benefit alone*
 - On the economic-growth effect: transitioning from a population growth rate of 3.6%pa to one of 2.5% pa over the period 2005-2050
 - Would cost \$25 per person in family planning service costs
 - Would increase per capita income by a (discounted) cumulative total of \$15,000
 - Benefit:cost ratio of 600:1 (!) from this benefit alone*
 - Benefits not counted:
 - Effects of slower population growth on political stability, climate change or environmental sustainability
 - Empowerment of women, gender equality
 - *Cost not counted(??): The loss of lives not born...*

*I have been unable to reproduce Kohler's calculations; the figures above are the ones I derive using Kohler's data and his stated methodology. The benefit-cost ratios Kohler reports are significantly different, but not in ways that affect the present discussion.

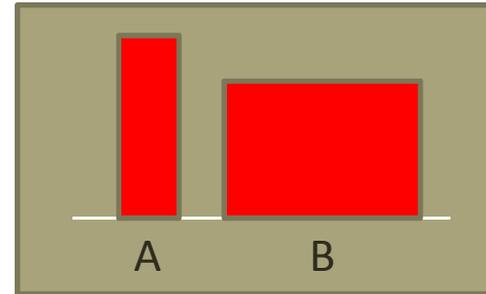
Population axiology: The ‘intuition of neutrality’, average and total utilitarianism

- Natural thought:
 - What’s bad is if *there are people* who lose most of their lives due to death in childhood. It’s not that ‘non-births’ are tragic, on account of the ‘loss’ of life that is not a loss that befalls any actual person. (Merely possible people don’t matter!)
 - This is because: adding an extra person is in itself *neutral*. (“The intuition of neutrality”)
- Broome (2004): This intuition is extremely natural, *but ultimately incoherent*.
- ‘The principle of equal existence’: Adding an extra person (others’ well-being levels being held fixed) leads to a state of affairs that is *equally as good as* the status quo.
- This principle is the relevant expression of the neutrality idea, but it cannot be true.



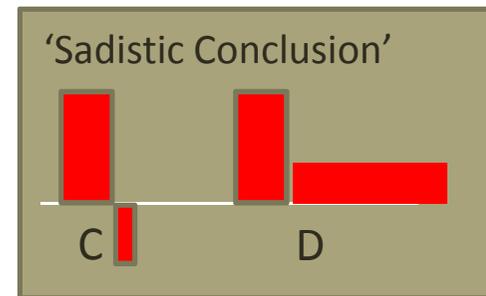
Beyond the intuition of neutrality: average and total utilitarianism

- The basic question of ‘population axiology’:
 - How should we compare states of affairs in terms of better and worse overall, when those states of affairs differ over the number of people who (ever) live?
- The two most-obvious candidate answers:
 - Average utilitarianism: The overall value of a state of affairs is the *average* well-being level in that state of affairs. (Therefore $A > B$.)
 - Total utilitarianism: The overall value of a state of affairs is the *total* amount of well-being in that state of affairs. (Therefore $B > A$.)
- Average, but not total, utilitarianism would (maybe) sanction ‘not counting the value of lives not born’.



Against average utilitarianism: The ‘Sadistic Conclusion’

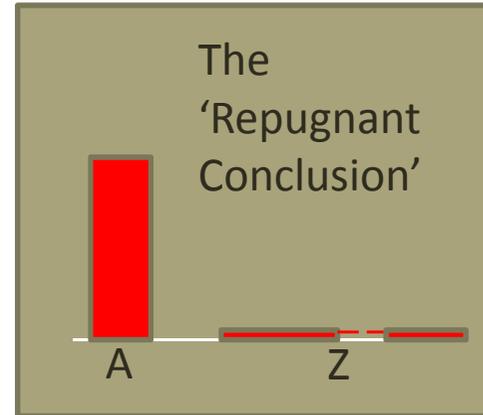
- The ‘Sadistic Conclusion’ (SC): It can be better to add some number of people with negative well-being, than to add (to the same ‘base’ population) a greater number of people all of whom have positive well-being ($C > D$).
 - The Sadistic Conclusion cannot be true. But average utilitarianism implies that it is. Therefore average utilitarianism is false.



Against total utilitarianism?

The 'Repugnant Conclusion'

- The Repugnant Conclusion (RC):
For any population A, there exists a strictly better population Z in which (however) no individual has a life that is more than 'barely worth living'.



- Total utilitarianism implies RC.
 - Repugnant?
 - But it turns out to be *very* difficult to come up with a theory that avoids RC, without running into (even) worse problems. (Arrhenius (ms))
 - And it doesn't follow that total utilitarianism recommends increasing the population until we reach the 'Malthusian limit' *in practice* (it won't).
- It is worth (at least) taking total utilitarianism seriously.

Sensitivity analysis for family planning calculations

- In general, cost-benefit analysts are careful to carry out ‘sensitivity analysis’ w.r.t. the assumptions needed for their calculations.
 - “One could... conclude that the task of estimating benefit-cost ratios is so difficult that it would be better to abandon it. But that would leave society with little systematic guidance about policy choices in this important area. Therefore... we swallow hard and proceed boldly ... to make the best estimates that we can given the present very imperfect information and strong assumptions necessary, with some efforts to explore the sensitivity of our estimates to important alternative assumptions.” (Kohler 2012, p.43-4; emphasis added)
- But, on population ethics in particular:
 - “Family planning programs, through their effect [on] fertility, affect the size of the population. This of course gives rise to the question of how to consider the welfare of persons who may not be born as a result of the intervention – a question that has been notoriously difficult to answer and for which no consensus exists in the literature. ... [W]e ... will not consider in the evaluation of family planning programs the welfare of individuals who are not born as part of the program.” (Kohler 2012, p. 43)
- This is unprincipled. We need to do sensitivity analysis w.r.t. population-ethical assumptions too, no less than other assumptions.

Family planning for total utilitarians (I): Life-years saved and lost

- In addition to averting 640,000 newborn deaths and 150,000 maternal deaths, Kohler's intervention scenario ('meeting all unmet need') also *averts 22 million live births* (annually).
 - If these 'lives averted' are valued (as 'lives lost') in the same way that Kohler's calculation values 'life lost' when a newborn infant dies, they amount to a *negative* effect of the family planning intervention, of magnitude \$1,430 billion.
 - Including this effect would transform the 'health' component of Kohler's benefit:cost ratio from +13:1 to minus 384:1.
- Two 'spins' on this (I advocate both):
 - Total utilitarianism is (probably) true, so Kohler's analysis is (probably) radically inaccurate.
 - Total utilitarianism is at least *a serious contender*, so Kohler's analysis dramatically fails the population-ethics 'sensitivity analysis'.

Family planning for total utilitarians (II): Economic progress

- Kohler's analysis assumes that the growth rate of *aggregate* GDP is unaffected by the population growth rate, so that growth rate of GDP *per capita* increases when the population growth rate decreases.
 - This is prima facie surprising – it requires a mutual 'cancellation' of numerous effects. ('Malthusian effect', 'Solow effect', various age structure effects, 'childcare effect', 'child quality effect', 'child quantity effect'...)
 - But (even assuming that this is correct), this is not obviously an *improvement*, unless neutrality or average utilitarianism is assumed: the increase in GDP *per capita* comes at the cost of many fewer people getting to live at all.
- Doublethink, part II: Ignoring the most *direct* effects of increased family planning

Rehabilitating family planning?

- The preceding analysis may be shortsighted: like the earlier ones, it altogether ignores the negative effects of rapid population growth on e.g. political stability, environmental sustainability, climate change...
- But
 - If this is where the real argument lies, then *we need to see the cost-benefit calculations* for these, harder to measure, effects.
 - There's no particular reason to think that family planning would then still top the global-prioritisation charts (many very good projects do not).
 - If (however) family planning *does* still top the charts, then the cost-benefit calculations for *mortality-reduction* projects are likely to be radically incomplete, and may (when completed) no longer be chart-toppers.
 - The basic issue is probably: **how do society-level neo-Malthusian concerns about overpopulation compare, in magnitude, to the intrinsic value of individual human lives?**

Summary and conclusions

- Child mortality and family planning are both (fairly) frequently cited as ‘top picks’ in global prioritisation.
- This is prima facie curious, since the most-obvious effect of the second intervention is precisely to undo the most-obvious effect of the first.
- Benefit-cost analyses (indeed) only manage to make both interventions simultaneously come out as ‘top picks’ by engaging in ‘doublethink’: making inconsistent decisions as to which effects (‘direct’ vs ‘indirect’) to count vs disregard, across the two interventions.
 - Analyses of mortality-reduction projects neglect indirect (e.g. economic) effects.
 - There *may* be a case for ignoring such effects in *some* contexts (e.g. doctor-patient relationships), but not at the level of global prioritisation.
 - Analyses of family planning programs ignore the (‘direct’) ‘value of lives not born’, counting *only* the ‘indirect’ effects on others.
 - This presupposes a person-affecting and/or an average-utilitarian approach to population ethics. Those approaches are initially intuitive, but ultimately indefensible.
- There is a resulting danger that we are currently wasting billions of dollars per year, by doing *and then undoing* good.
- To fix this: More sophisticated analysis, *including serious attempts to put neo-Malthusianism and the value of individual additional lives in dialogue with one another*, is required.

References

- Arrhenius (ms), *Population ethics: The challenge of future generations*
- Brock (2003), “Separate spheres and indirect benefits”
- Broome (2004), *Weighing lives*
- GiveWell, <http://www.givewell.org/international/top-charities/AMF#Whatdoyougetforyourdollar>
- Kohler (2012), “Challenge paper: Population”
- Lippert-Rasmussen and Lauridsen (2010), “Justice and the allocation of healthcare resources: Should indirect, non-health benefits count?”