Intro

I have drafted two entries into the Legal Priorities Project Writing Competition on “Improving Cost-Benefit Analysis to Account for Existential and Catastrophic Risks”.

I am not a US citizen so cannot enter the competition.

I am interested in having these ideas presented and possibly incorporated into LPP advocacy. Mackenzie Arnold has advised that this is possible and that these ideas can be presented to the nominating panel for them to provide feedback and if they find them promising, passed to the final round judging panel for feedback as well.

Some of the notes below are rough, I have also not invested as much time into this as if I was able to enter the competition. I am also happy if others take and use these ideas in order to enter the competition.

Contents

The two ideas were

1. Going beyond quantitative analysis

Summary 1

Circular A-4 provides advice to regulatory analysts. However it focuses almost exclusively on quantitative Benefit-Cost Analysis (BCA). This risks missing the woods for all the trees. There are clear limits to BCA and it is important to only use BCA in conjunction with other kinds of methodologies. There are many other types of evidence out there that decision makers can be made aware of that analysis should look into. There is good precedent for this, for example advice for government analysts in the UK suggests they investigate a range of considerations for supporting decision making above just quantitative models.

Suggested amendment 1

I suggest radically amending the kind of analysis that Circular A-4 guides analysts provide to decision makers along the following lines.

Amendments to the introduction should be made as follows:

> A good regulatory analysis should include the following ~~three~~ seven basic elements: (1) a statement of the need for the proposed action, (2) an examination of alternative approaches, and (3) an evaluation of the benefits and costs—quantitative and qualitative—of the proposed action and the main alternatives identified by the analysis (4) an evaluation of the past success, case studies, past precedent and quality of evidence for the proposed action (5) an evaluation of the strategic and long-term case (6) an evaluation of practical considerations for delivering the proposed action, and (7) a summary of expert views on the proposed action.

Furthermore additional sections (drafted in short here, additional elaboration would likely be needed) should be added to the guidance after Section G to cover:

> **H. Case study and historical analysis**
>
> Analysts should look for cases where the US government or other governments or large institutions have had to solve this specific problem in the past. They should consider what actions were taken in that case, how well those actions worked and how relevant that case is to the current situation. They should also consider case studies of analogous interventions, when similar problems were being addressed or similar interventions were being tried for different problems. They should actively look for scenarios where interventions were unsuccessful and remember there will be publication bias toward those cases that did work. This should be summarised in roughly half a page setting out to decision makers what if any existing best practice there is for the current policy decision and additionally for each option a summary of how policy makers adopting that style of approach fared in past scenarios.

I. Analysis of strategic and long-term considerations

Analysts should investigate and then summarise for decision makers key theoretic, strategic and long-term considerations. This should be summarised into roughly one page. For example the summary could take the form of a table using red amber green ratings for each strategic criteria for each option alongside additional key details. This note should include:

A strategic background: Why this situation is considered a market failure requiring government intervention and why it has not been successfully addressed to date. This should be based on economic and historical analysis.

Long-term considerations: The extent to which each option moves the country towards a shared vision of a good future. This should be based on a broader exercise that the US government should carry out to carry out public consultation and map out a long-term vision for the country so as to support long-term planning.

Robust to future scenarios: How each option fares under a range of four key future scenarios. This should be based of an analyst driven scenario planning exercise

Robust to future trends: What are the most relevant long-term future trends and how this affects each option. This should be based on a trends analysis exercise.

Robust to risks: Do any options break down if the county was undergoing a crisis situation. This should be based on a broad national risk assessment exercise and considering how the options would fare under a national disaster.

Any other strategic considerations the analysts deem valuable

J. Analysis of practical considerations

Analysts should investigate and then summarise for decision makers key practical considerations. This should be summarised into roughly one page. For example the summary could take the form of a table using red amber green ratings for each strategic criteria for each option alongside additional key details. This note should include:

Practical considerations: How deliverable is each option. What are the key barriers to delivery such as: enforcement, procurement, financing, management, skills and staffing, public perception, and political opposition. This should be based on mapping out the steps of theory of change and investigating each step through engaging relevant external stakeholders.

Simplicity assessment: How many key steps are in each option (more simple interventions being more likely to work)

Likely failure modes: For each option the top 3-5 ways it could fail to deliver as expected should be listed. This should be based on a theory of change assumption mapping exercise

Any other practical considerations the analysts deem valuable

K. Summary of expert views

A table summarising the views of each expert or key group of experts, and their potential biases or vested interests should be drawn up to support decision makers. This should be one page or less.

Essay 1

What are quantitative models?

A quantitative model analysis is a form of analysis most commonly used in economics, health economics, and evidence based policy making. They consist of one or more calculations and result in a numerical output. There are a variety of different types of quantitative models used. The most common of which are a Benefit-Cost Analysis (BCA) and a Cost-Effectiveness Analysis (CEA). A BCA consists of one or more calculations and results in a ratio of the costs of a given action or intervention relative to its impacts with costs and benefits generally measured in dollars. A CEA will compare the costs to the main impact where the impact is measured in its own units, such as Quality Adjusted Life Years or lives saved.

Note on terminology: as Circular A-4 focuses on BCA the rest of this essay also focuses on BCA, although the points made also apply to CEA and any other quantitative models.

These quantitative models output a single (or small set of) endline numbers that represent the value of different policy options. Options that have a higher ratio of benefits to costs are considered better than options with a lower benefit cost ratio (BCR), all else being equal. This is useful for deciding between options or between taking action and continuing business as usual.

There are many benefits to using such quantitative models to guide decision making. The key benefits are listed here (in a rough order of strongest to weakest):

- Clearly connect to endline goals
- Can be used to compare interventions that are otherwise difficult to compare
- A clear and simple metric allows for transparency, accountability and oversight
- Reduce biases and discourages decision making on simple ideological heuristics
- Allow for formal sensitivity analysis
- Are a respected tool in multiple fields
- Considers scope
- Can lead to novel conclusions
- Encourage quantitative analysis more broadly

The limits of quantitative models

BCA has limitations.

It is important to distinguish between the "true benefit cost ratio" of an action and the "modelled benefit cost ratio". The true BCR of an action—if it's known—would be a highly relevant metric and could be weighted very heavily when making a decision. However, the closest an evaluation-focused organisation can usually get to ascertaining the true BCR of an intervention is through constructing a model, which is almost by definition an imperfect

estimation. This is because we often lack important data about the world, or a sufficient amount of it. As such organisations will end up with a modelled BCR or a modelled expected value of an action. This needs to be considered in light of the limits of a quantitative modelling approach.

Flaws of BCAs listed here from strongest to weakest:

- BCAs simplify reality
- Subject to the “optimizer's curse”
- Necessarily involve value judgments
- Prone to mistakes
- May not be generalizable to other contexts
- Make it hard to model flow-through effects
- Can be misleading in many ways
- The interventions we analyze are somewhat preselected for cost-effectiveness
- Subject to researcher bias
- May bias you towards interventions with more measurable results
- 90% confidence intervals can be misleading
- Are often not adjusted for priors

BCAs simplify reality: As mentioned above, cost-effectiveness models are necessarily simplifications of reality, not a perfect representation of reality. This is both a strength and a weakness. Although it allows us to get a clearer understanding faster, it also means that they do not accurately capture reality. Adjustments in the variables used will change the final value of the BCA.

BCAs are subject to the “optimizer's curse”: All estimates are prone to error, and these errors compound. An intervention whose BCA yields a high cost-effectiveness is more likely to have had errors in its favor. This means that the most and least cost-effective interventions are likely to regress to the mean upon further examination. Overweighting BCAs in our decision making could lead us to neglect good opportunities that did not have as many favorable errors. This is less of a problem in richer information environments.

BCAs necessarily involve value judgments: It is surprising how much value judgments can differ. For example, a simple cost effectiveness model in an area with a lot of data, such as health economics, will need to make assumptions such as comparing the "value of averting a death" compared to the value of "doubling consumption for one person for one year". These numbers might vary depending on the life circumstances of the affected individuals in question, the average age of the affected individuals and so on. Ultimately this will rely and value judgments and reasonable estimates could be an order of magnitude different from one another. If these value judgments are subjective preferences that vary among individuals, then BCAs are only generalizable insofar as the researcher’s values align with the reader’s.

BCAs are prone to mistakes: Mistakes are inevitable, due to human error and/or poor information quality. Although small mistakes usually only translate to small problems on their own, these mistakes compound in a multivariate model, thus exaggerating the consequences. Independent analysis of quantitative models by external organisations often highlight numerous such mistakes. For example, an analysis by charity evaluator GiveWell

into the Disease Control Priorities Project estimates for the impact of deworming found 5 mistakes that contributed to an overestimation of the intervention's cost-effectiveness by one hundred times (source).

BCAs may not be generalizable to other contexts: Some BCAs rely heavily on randomized controlled trials (RCTs) for their data, and in some cases, this can be problematic. If an RCT was conducted in one particular region or with one particular method, the effect size may change dramatically in different regions or with other methods.

BCAs make it hard to model flow-through effects: Researchers have written that it is difficult to properly model flow-through effects in BCAs. Indeed, a common tactic is to ignore flow-through effects entirely. There are solutions to this problem; however, they all take vast amounts of time or are prone to error.

BCAs can be misleading in many ways: If researchers fail to consider important factors or are not transparent in their reasoning, BCAs can yield misleading results. For example, if a BCA concerns an expected value, the probability of success must be clarified. If only pure expected value is reported, there is no difference between a 50% chance of saving 10 children and a 100% chance of saving five children. This would fail to consider any level of risk aversion.

The interventions we analyze are somewhat preselected for cost-effectiveness: As the sources for policy ideas are not random these ideas will have been created with cost-effectiveness in mind. Given the "optimisers curse" this means that random error will account for more of the variance, making BCAs a weaker tool.

BCAs are subject to researcher bias: BCAs are resistant to certain biases, but are susceptible to others. If the researcher conducting a particular BCA has a favorable view of the intervention, for example, he or she may (consciously or unconsciously) bias the results in its favor. A researcher's desire to find novel, cost-effective interventions may also have this result.

BCAs may bias you towards interventions with more measurable results: Effects that are difficult to measure may increase the error rate or be neglected. This can lead to an underestimation of the effectiveness of interventions with hard-to-measure outcomes.

Ninety-percent confidence intervals can be misleading: Depending on how well calibrated researchers are, the worst-case scenario, the best case and 90% confidence interval (CI) may be incorrect. CIs are particularly susceptible to this, as we are likely to underestimate the range of uncertainty that is actually accurate. Worst case and best case are no better, as these may rely on many unlikely events happening, meaning the probability of either occurring is minimal.

Are often not adjusted for priors. As per Beysian statistics quantitative estimates should be adjusted using a 'Bayesian prior'. It is hard to incorporate prior beliefs of relevant actors into a model, especially if a model is being created to allow a group of relevant actors with varying priors to discuss and make a decision. In these cases it can often lead to better

decision making for Beysian adjustments to be left out of the model and to come out of the discussion.

Cannot capture unknown unknows. Quantitative models work where data is available. However in scenarios that involve risks and unknown unknown factors that could affect a policy outcome it can be almost impossible to model this quantitatively.

Concerns with reliance on BCAs in evaluation of projects have been discussed in depth elsewhere. The theoretical concerns are outlined by GiveWell, 2011 and the comprehensive coverage of the practical concerns are outlined by Saulius Šimčikas, 2019.

How much weight should be given to quantitative models?

It is clear that given the limits of BCAs, so what is the solution?

A key solution that can work to address many of these limits is to carry out additional analysis.

An analyst team could carry out multiple quantitative analyses, done independently by different researchers with different values and biases, using different methodologies, and presenting results in different formats. If these multiple models converge this would add certainty to the decision making and if they diverge an additional investigation could be carried out.

In scenarios where quantification is easy and data is readily available, multiple quantitative models might be the best approach. However in more high uncertainty situations, those with limited data and unknown unknowns it is often better to use a broader variety of tools to make a decision than just quantitative modelling. Most regulatory policy decisions likely fall under this category. There are many other types of evidence out there that decision makers can be made aware of that analysis should look into.

Focusing solely on benefit cost models is like missing the woods for all the trees. So much evidence needs to be collected to do a good benefit cost model but that evidence can be more directly used by decision makers. For example Circular A-4 suggests trying to identify uncertain numbers in a CBA by engaging experts and using tools such as the Delphi methods:
“ expert solicitation is a useful way to fill key gaps in your ability to assess uncertainty. … These solicitations, combined with other sources of data, can be combined in Monte Carlo simulations to derive a probability distribution of benefits and costs”.
However, expert solicitation and Delphi methods can be used directly to ascertain expert views on the key policy options under consideration, not just as part of a process to develop a quantitative model to decide on policy options. A decision maker given the outputs of a quantitative model will make a worse decision than one given the outputs of a quantitative model and a short summary of the opinions of relevant experts and details of who those

experts are. Providing a short summary of expert views should be a part of the job of regulatory analysis.

There are good precedents for this type of approach, for example from the UK.

Case study: the UK government's use of quantitative models?

The UK Government Green Book sets out the process to be used in analysis by UK regulatory decision makers. In this process the benefit-cost is just one of many factors that a government analyst needs to look into and provide advice on for decision makers.

Box 5. The Five Case Model

Strategic dimension	**What is the case for change, including the rationale for intervention?** What is the current situation? What is to be done? What outcomes are expected? How do these fit with wider government policies and objectives?
Economic dimension	**What is the net value to society (the social value) of the intervention compared to continuing with Business As Usual?** What are the risks and their costs, and how are they best managed? Which option reflects the optimal net value to society?
Commercial dimension	**Can a realistic and credible commercial deal be struck?** Who will manage which risks?
Financial dimension	**What is the impact of the proposal on the public sector budget in terms of the total cost of both capital and revenue?**
Management dimension	**Are there realistic and robust delivery plans**? How can the proposal be delivered?

As well as the economic benefits and financial costs, UK government analysts also investigate both the strategic case and the practical (commercial and managerial) factors for any regulatory policy option.

Quantitative analysis is of course a key part of the final stage analysis but only applied after all options have been assessed for their strategic fit and practical deliverability.

Choice of additional tools for regulatory analysts

[Additional notes could be added on what the best tools for policy makers to use are on top of BCA, and why it should be the role of analyst teams to research these tools.]

Conclusion

Multiple sources of evidence should be investigated by analysts and provided to decision makers alongside the outputs of any BCA.

Note: significant amounts of this text are borrowed with permission from https://www.charityentrepreneurship.com/cea

2. Adjusting for presentism and optimism bias

Summary 2

Quantitative analysis is prone to be affected by analysts and decision makers' biases. Explicitly adjusting for these biases should be done as part of the analysis. There is good precedent for this, for example advice for government analysts in the UK suggests this. Adjustments can be made for both optimism bias (especially for any decisions that require government procurement or major infrastructure to be developed) and for presentism bias.

Suggested amendment 2

In section H the following should be added

> **Adjustments to account for biases.**
>
> When conducting appraisal consideration should also be given to optimism bias – this is the proven tendency for appraisers to be optimistically biased about key project parameters, including capital costs and operating costs, project duration, and resulting benefits delivery. Optimistic rather than realistic projections result in undeliverable targets and if permitted across the board create institutional failure as all proposals fall consistently far short of promised results. For this reason, specific optimism bias adjustments must be applied as numbers are initially identified. Ideally adjustments should be based on an organisation's own evidence base for historic levels of optimism bias. In the absence of robust organisation-specific estimates generic values are provided in the table below. There are currently no generic values available to be applied to benefits, only to costs, however an adjustment should be applied based on an organisation's own evidence base.
>
> *[A table should then be inserted with figures for different kinds of cost inputs and how to adjust them for optimism bias, based on historical analysis of past BCA]*
>
> When conducting appraisal consideration should also be given to presentism bias – this is the tendency for decision makers to favour the short run benefits over the long run benefits and disfavour short-run costs over long-run costs. This leads to short term decision making. For this reason, specific presentism bias adjustments can be applied at the same time as a discount rate is applied. This should be in the form of a small positive amount that effectively lowers the discount rate. Ideally adjustments should be based on an organisation's own evidence base for historic levels of presentism bias. In the absence of robust organisation-specific estimates a generic value of 1.5% per annum should be used. The resultant output should be provided to decision makers alongside the non-adjusted figure to support their decision making (not as a replacement to the unadjusted output)

Note text borrowed from the UK government Green Book.

Essay 2

Notes:

- There is a precedent for adjusting for biases as part of quantitative analysis, for example from UK government Green Book sets out adjustments for optimism bias.
- Adjusting for presentism bias and offering decision makers a prentism bias adjusted output figure, alongside the standard figure, would help ensure that policy makers have the option to choose something that is good for the long term (e.g. investing in pandemic preparedness).
- Note: There is no strict universal definition of “presentism bias”. For the purpose of this suggestion we are referring to the tendency for decision makers to favour the short run benefits over the long run benefits and disfavour short-run costs over long-run costs. This is also known as political short-termism.
- There is various evidence for political short-termism and the need to address it, for example see papers here and here and here and books such as this and this and this.
- A case could be made to US policy makers for addressing political short termism on the grounds that it is less of a challenge to decision making in autocratic regions (such as China). Although democratic regimes (such as the US) clearly lead to better decision making in many cases, long-term decision making does suffer. If we want to US to be globally competitive in the long-run we need to find ways to support long-run decision making.

[Additional details to be added]