tailrisk economics
te maha hoki o nga pihikete

The mortality reduction benefits of insulation: the error identified

Our paper showing a critical error in an MBIE cost-benefit analysis that grossly over-stated the benefit of insulation has raised something of a stir.

We claimed that there was double counting of the mortality benefits and suggested that this may have been due to the wrong number, the 5-year present value rather than the annual benefit, being picked up for the 30-year present value calculation.

The authors of the underlying cost benefit study have come back and said that they picked up the 'right' number and it was us who have made the mistake.

We were correct about the double counting. However, we were wrong about the nature of the mistake. It was not an inadvertant pick-up of the wrong number. There was a fundamental logical error in the modelling.

The analysis assumed that there was not one, but five occupants who had been hospitalised with a cardiovascular illness in the previous 18 months in each of the relevant insulated houses. There should have been only one such occupant.

Our criticism stands, a five year present value of the benefits has been used as a measure of the annual benefit, rather than the actual average annual benefit.

To explain the error we will walk through the following steps:

1. A review of what the authors said that identifies the logical error
2. A walk-through of a concrete example of the benefits from an insulated house.

3. A ‘top down’ sense test of the aggregate numbers.

1. What was said

The authors’ thinking is captured in the following passage from page 38 of the report.

“*We have modelled the probability of a vulnerable person avoiding mortality as a result of the intervention. The probability of this is (112.7/1000)*0.27= 0.03 (3%). We treat avoidance of mortality by treatment in each year as independent events. The multi-year benefit calculated above would accrue based on the life years gained as a result of deaths avoided in year one.* ***However, we would expect these benefits to accrue in year two for different vulnerable individuals (aged 65 and over with a cardiovascular related hospitalisation in previous 18 months), and for different individuals again in every subsequent year that the treatment continues to have an effect, i.e. an on-going stream of benefits of $1,050.74 per year. This assumes a constant proportion of people aged 65+ who have recently been hospitalised with circulatory problems…..***

If we estimate that a person who does not die as a result of receiving treatment lives for half as long as they would otherwise (conservatively 5 years) then we will end up with the following on-going annual saving per dwelling (discounted at 5%):”

$$= \$63.96 + \sum_{n=1}^{n=4} (\$127.92)/((1+0.05)^\wedge n)$$
$$= \$517.59 (95\%CI : \$0.00, \$900.99)$$

The section in bold gives a clue what is going on here.

The authors would have wrestled with the modelling of benefits with dwelling and occupant populations that had a static average composition, but with aging individuals. They have assumed that the original occupant will be replaced by a different individual in the second year, the second with a third in the third year and so on.

The logic appears to be that as each individual will be in their own first year of occupancy it is appropriate to add the discounted ‘first year’ benefits to get the ‘annual’ benefit.

The logical error here is that the first occupant will still be living in the house (the annual mortality rate is about 10 percent) when the second occupant arrives. So the authors have inadvertently increased the occupancy in the house from one to five by the end of the five years, and so have amplified the ‘annual’ benefit.

It doesn't actually matter how long the occupants live, as long as, statistically, they are replaced by a similarly vulnerable occupant. And as there can only be one occupant at any point in time the benefits **must** be calculated as the present value of all of the annual benefits (the $128 average per household) over the 30-year horizon.

What obviously doesn't work is the implicit assumption that the representative house becomes a kind of elderly 'Young Ones' flatting arrangement. We can imagine a scene where an elderly, and very crotchety, Vyvyan (Adrian Edmondson) tells Rick (Rik Mayall) "Piss off, you are not moving into my nice insulated house. I am not bloody dead yet you bastard. The modellers gave me five years".

Another way of expressing the basic point here is that an annual benefit is just that, the benefits that accrue in a particular year, it does not include benefits accruing in future years regardless of whether they accrue to a new occupant or not.

To sum up, the authors have calculated a five-year present value, but have misinterpreted it as an annual figure through a misunderstanding of the logic of their model. They implicitly put five elderly occupants into a dwelling that should only have one.

We must admit that we got the source of the error wrong in our document. The correct annual benefit number seemed so obvious that it did not occur to us that the methodology being described was to multiply the number of occupants by a factor of five. We though the most likely error was that the wrong number was picked up.

Modelling how aging individuals interact with a static population of houses and occupants is complicated, and the authors simply made a mistake. The mistake is obvious once it is identified, but it would not have been as obvious at the time. However, a 'top-down' sense test along the lines of the one presented in section 3 below, by at least one of the chain of analysts who have reviewed or used the model outputs, would have revealed that something was awry.

2. A more concrete example

To understand what is going on here more clearly, we found it useful to work through a simple example starting with a single insulated house rather than work with statistical averages as the authors have done. All of the inputs into our estimates are from the authors' papers.

We assume that the initial 77-year-old occupant who insulates the house will live for exactly the average 5 years, and then will be replaced by another 77 year old and so on for the 30 years. There are six successive occupants.

- The reduction in morality is 27 percent, so insulation will reduce the occupants' mortality rates by 3 percentage points.
- The value of a life year is $150,000, so the expected annual value of the mortality reduction is $150,000 x .0304 = $4560
- The present value of the stream of benefits over 30 years with a discount rate of 5 percent is $4560 x 15.37[1] = $70,098
- The authors apply a 15 percent reduction to benefits for technical reasons so applying the same adjustment gives an expected present value per dwelling of $59,584.

Our reading of the cost-benefit paper was that the stock of insulated dwellings was 178,000. As 2.8% are occupied by the vulnerable elderly, the number of relevant dwellings is 4984.

So the total NPV of the mortality reduction benefit is:
4984 x $59584 = **$297 million.**

3. A top down calculation

The top-down calculation works as follows:

- The number of lives saved per year per 1000 dwellings is 0.852.
- A total of 178,000 treated dwellings gives 151.7 lives saved per year.
- At $150,000 per life year the annual benefit is 151.7 x $150,000 = $22.7 million.
- The present value of a stream of benefits of $22.7-million over 30 years with a 5 percent discount rate (after applying the technical adjustment) is **$297** million.

The total present value benefit estimate reported in the cost-benefit paper is $1263million (page 22). A sum of $923 million relates to the lower morality benefits.[2] This figure is clearly wrong. It cannot be derived using the inputs from the study. It was generated by the error in calculating the annual benefit. Using the inputs from the study we calculated the benefits to be about $300 million.

---

[1] For simplicity the PV multiplier assumes benefits accrue at the end of the year. This slightly understates the benefits.

[2] Note that it appears that a discount rate of 4 percent is used in this paper rather than the 5 percent in the sister paper where the methodology was explained.

Contact: Ian Harrison – Principal Tailrisk Economics

ian@tailrisk.co.nz

Ph. 022 175 3669

04 384 8570