Review of Shadbolt et al, Methane Science and Target Review

Myles Allen

I am Professor of Geosystem Science and Head of Atmospheric, Oceanic and Planetary Physics at the University of Oxford, submitting this review in my personal capacity. I have published extensively on the challenge of how to compare methane and carbon dioxide emissions. I was one of the authors of the FaIR model used by the Review, although I did not contribute to the representation of methane lifetime and am not an expert in atmospheric methane chemistry – this is important, because most of my concerns are on this latter topic. I have visited New Zealand twice in recent years to lecture on atmospheric science and climate policy. Regarding potential conflicts, I supervise a doctoral student whose studentship is supported by Hilton Food Group (subject to Oxford's strict standards of researcher independence), and junior members of my research group (although not I myself) received funding to conduct and write up a short modelling study commissioned by Beef and Lamb New Zealand (Barth et al, 2023).

Overall, the Review is commendably workmanlike, concise and to the point. It answers the question as set, while correctly identifying the key ambiguity, which is that future methane emissions consistent with no further methane-induced warming depend on mitigation policy decisions outside of Aotearoa New Zealand (henceforth NZ)'s control. These issues are detailed further in the section on global policy ambition and the warming impact of NZ's emissions, below. First, however, I have some comments on the overall framing.

Overall framing of the Review

It would be helpful if the Review could make it even clearer exactly what the question is that it is answering, and also (if it isn't out of scope), reminding people of the questions it is not answering. The question is specifically what rates of reduction of methane emissions are consistent with returning the global warming caused directly by NZ's methane emissions to its 2017 level by 2050 or 2100. I'm not sure if it was a Review decision to interpret "no further warming" in this way, or whether it was specified in some accompanying guidance (the ToR is not so specific), but this needs to be spelled out, because the casual reader would probably not realise that, under this definition, warming induced by NZ's methane emissions can rise above 2017 levels before falling back again (as, indeed, it does, in all scenarios).

I think it would be helpful to emphasise (and perhaps give indicative numbers) that this is additional warming above a substantial (in per capita terms) warming prior to 2017 due to historical NZ emissions, primarily methane since the 19th century, and deforestation before then. It would also be helpful to make clear exactly what is meant by the direct warming impact of NZ's methane emissions, viz. how global temperatures would be different if the entire impact of NZ's methane emissions, including prior to 2017, were absent. The impact of eliminating NZ's methane emissions in 2017 or 2025 (supposing that were possible) would be considerably greater than this additional warming, and hence there is a sense in which any ongoing NZ methane emissions indirectly cause warming as long as global methane-induced warming continues to increase, because they contribute to a global pool

of methane emissions that is driving up global temperatures. Acknowledging these different possible interpretations of "causing warming" would be helpful, but I appreciate this may be out-of-scope.

Global policy ambition and the warming impact of Aotearoa New Zealand's emissions

There are two contributors to this dependency: the impact of atmospheric composition on (1) methane radiative efficiency, or the radiative forcing per unit additional atmospheric methane loading; and (2) methane oxidation rate, or the fraction of methane in the atmosphere that is oxidised every year, generally expressed as a lifetime (the reciprocal of the oxidation rate). Together, these determine the impact of NZ's future methane emissions on the future global energy budget and hence global temperature. Broadly speaking, the more ambitious climate policy is in the rest of the world, the faster NZ's methane emissions need to decline to be consistent with causing no further warming.

The quantification of this dependence is where I have the most issues with the Review – to the point that I became concerned that I didn't understand what was going on and so repeated key calculations using an entirely independent (publicly available) model. To my great relief, I got almost exactly the same answer to the core question. So, the most important overarching point is that, although I do think it is worth clearing these issues up, I don't think they impact the Review's conclusions.

The key modelling tool used in this Review is the FaIR model, but the Review does not (as far as I can tell – apologies if I missed it) specify which version is used. This is important, because while FaIR2.0 (Leach et al, 2021) is the simplest version of the model supported by a comprehensive peer-reviewed publication (and hence the version I would still recommend to be used for a study such as this one), FaIR2.1 and FaIR2.2 (https://docs.fairmodel.net) offer a more comprehensive representation of methane chemistry, which affects methane lifetime.

I am confident of the treatment of effect (1), methane radiative efficiency in FaIR, because it has been extensively tested against the parameterisation of Etminan et al (2016) and has been stable, with only minor updates, in the FaIR model since its introduction in FaIR2.0. Fortunately, effect (1) is, I think (but would welcome the panel confirming), the stronger of the two contributors affecting how the warming impact of NZ's emissions depend on global mitigation ambition.

I am less confident in the treatment of (2), the methane oxidation rate or lifetime. In FaIR2.0 this was simply parameterised in terms of a dependence on atmospheric methane concentration and tropospheric temperature. More methane implies a lower concentration of OH radicals (because they are reduced by methane), and hence a longer methane lifetime, while higher temperatures imply faster oxidation rates for the same OH and CH4 concentrations, slightly reducing methane lifetime. Together, these imply a more-or-less monotonic increase in methane lifetime from 8.3 years in 1850 to 9.8 years in 2100 under a high emissions scenario, according to figure S2 of Leach et al (2021).

Figure 13 of the Review suggests a more complicated behaviour is displayed by the FaIR version used, with a methane lifetime gradually increasing to 1950 or so, followed by a rapid decline to the present, and then a broad range of lifetimes in future depending on scenario. This plot suggests to me that the parameterisation of methane lifetime used has been updated from FaIR2.0, and now includes dependence on other atmospheric constituents. What worries me is that Leach et al (2021) concluded “a key message to emphasize is that these non-linearities are not very significant on the whole, and especially on longer (centennial) timescales” which doesn’t seem consistent with the very different behaviour of figure 13 and figure S2 of Leach et al (2021).

To compound my confusion (and I really want to stress that I’m not an expert in atmospheric chemistry, so I may well be simply missing something obvious), the Review states after Figure 13 “if the world follows a high emissions and warming scenario such as SSP5-8.5, then methane lifetimes would reduce (due to the enhanced production of the hydroxyl radical).” It is not obvious to me that higher emissions automatically imply more OH radicals: simply adding more methane to the atmosphere would imply a more reducing atmosphere and lower OH. This seems to be supported by a quick survey of the literature: for example, Lui et al (2024), https://www.nature.com/articles/s41467-024-47436-9 state "if following a largely unmitigated scenario, the global OH exhibits a significant decrease that would exacerbate methane’s radiative forcing" through a longer lifetime, consistent with Leach et al (2021).

At a bare minimum, the Review must state the version of the FaIR model used and give the source of the code (downloaded from where, when). Given I don’t have specific expertise in this area, figure 13 should be checked with an expert in methane chemistry to see if it makes sense: I stress it may well do, methane chemistry is complicated, but it just doesn’t seem to be consistent with what I would expect. Then it needs to be explained in terms that I, as a non-chemist, would understand. Right now, the Review is vulnerable to someone claiming “their model gets the behaviour of methane lifetime wrong so we can’t believe anything it says.” In fact, it makes effectively no difference, as I show below.

Note: Since submitting my provisional draft review, I have been in communication with the Review authors, and appreciate their treatment of methane lifetimes is well within the range of uncertainty – but given this uncertainty is large, it is important to re-emphasise that it doesn’t matter.

Verifying key findings with the AR5 linear impulse response model

Concerned about the potential impact of the methane chemistry modelling, and appreciating the considerable time pressure on this Review and likely scrutiny of key conclusions, I thought it would be helpful to reproduce them using an independent model. For maximum traceability, I used the Impulse Response model that was developed for the IPCC 5th Assessment Report. Although this is a completely linear model, and hence less realistic than FaIR for modelling the total response to a finite-amplitude emission, it is very widely used to model the impact of small emission perturbations, including by governments and commercial entities to support mitigation policy and greenhouse gas trading, because it remains the model most widely used to derive values for Global Warming Potentials (GWPs) and Global Temperature-change Potentials (GTPs). Hence, despite its simplicity, this model

has the advantage that it is already endorsed for policy use, including by the NZ Climate Change Commission, who regularly use GWPs based on this model in the calculation of NZ's aggregate emissions. The response of this model turns out to be well within the range of responses of the non-linear model.

The only modifications made to the AR5 model were to adopt a value of the equilibrium climate sensitivity that is closer to the current best-estimate (3°C instead of 3.9°C in the original) and to use updated ocean thermal adjustment time-constants (Geoffroy et al, 2013). I have provided the code used, reproduced from Allen et al (2022), which has the advantage of a large number of co-authors from across the spectrum of views on greenhouse gas metrics.

Calculating the warming response in any given year to past emissions with this model is simply a matter of multiply the emissions in each preceding year by the methane AGTP value corresponding to how long ago that preceding year was and adding up the results:

$$T_{CH4}(t) = \sum_{k=1}^{t} E_{CH4}(t-k)\ \mathrm{AGTP}_{CH4}(k)$$

(assuming there were no emissions before year 0).

The impulse response functions on which the AR5 model is based linearise the response of the climate system to emissions of greenhouse gases against a background of constant atmospheric composition. While this is clearly unrealistic for carbon dioxide, it is at least within the range of future scenarios for methane, between the very substantial increases in methane concentrations in high-emission scenarios and equally substantial reductions in methane concentrations in high-ambition mitigation scenarios.

I would argue that linearising the climate response against constant global methane concentrations is at least as defensible as looking for commonalities in responses across scenarios as in Table 2. In effect, linearisation about a mid-range scenario is, to first order, equivalent to averaging the responses of a non-linear model across the range of scenarios, as the reader is implicitly invited to do in interpreting Table 2.

Results of a linear model are, by definition, not sensitive to the background scenario, and are shown in the figure. The left panel shows NZ emissions of the three major greenhouse gases, expressed as CO_2-e_{100} by scaling by their AR5 GWP_{100} values. Historical emissions (merging Primap and NZ inventory data, scaling the former to the latter) are shown to 2022, followed by a linear decline to zero in 2050 for CO_2 and nitrous oxide, and a 4% per annum decline in methane emissions from 2022 onwards, corresponding to a 10.7% reduction by 2050. This scenario returns NZ's total methane-induced warming in 2050 to almost exactly its 2017 level, consistent with the results shown in Table 2. Hence the conclusion that about a 10% reduction in methane emissions between 2022 and 2050 is sufficient to return methane-induced warming to its 2017 level is not dependent on the model used or the parameterisation of methane lifetime.

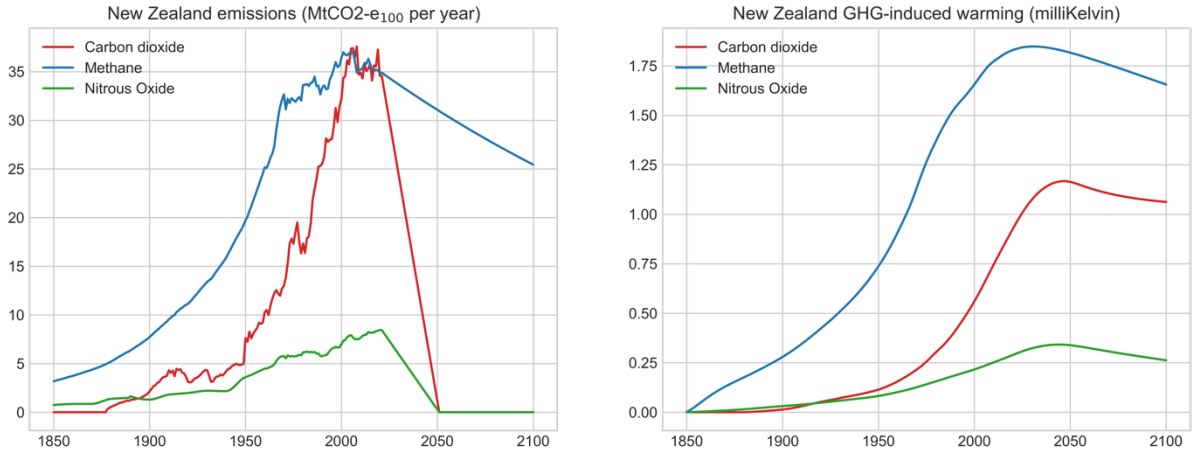

Figure: Left panel: NZ's emissions of the three main GHGs, expressed as CO_2-equivalent using the conventional GWP_{100} and, Right panel: warming response calculated using the AR5 linear impulse response model. From 2022 CO_2 and nitrous oxide decline linearly to zero in 2050, while methane declines by 0.4% per year. Calculations using code provided with Allen et al (2022).

While this figure is provided simply for the information of the panel, I do urge them to consider using something like it (and provide all the code for their convenience), because it makes the point that the overall shape of NZ's CO_2-induced, nitrous-oxide-induced and methane-induced warming are similar, with the key difference being that methane-induced warming is stabilized rather earlier than the other two.

Hence my key substantive asks of the Panel are:

a) Document the origin and version number of the FaIR model used.
b) Explain the variation of methane lifetime shown in 13 and justify it with reference to the behaviour of explicit atmospheric chemistry models (it looks to me to be inconsistent with the results of more complex models, but there may well be a good explanation).
c) If (b) proves impossible within the time-frame, re-do calculations with a version of FaIR with a comprehensible and traceable methane lifetime parameterisation, such as FaIR2.0, and consider showing results from the AR5 linear response model in an appendix for confirmation.
d) In the unlikely case of both (b) and (c) proving impossible, consider presenting core results with the fully-traceable AR5 model, updated to AR6 parameters if that is deemed a higher priority than traceability to the model currently still the most used for GWP calculations.

Not wishing to lose another day with the time-zone changes, I am sending this as an interim report, along with the data and Python notebook. I will also send an annotated copy of the PDF with presentational suggestions, but this is the substantive one.

I hope this interim review is useful, and would be happy to communicate directly with the panel members if anything needs further clarification.

Myles Allen