Responses to CASAC Questions on the Ozone ISA from Consultant Dr. Lorenz Rhomberg

Questions from Dr. Cox

Dr. Cox's set of questions parallels to a large extent the set he posed in regard to the $PM_{2.5}$ ISA. To a large extent, the questions are posed as methodological quandaries rather than as questions about specific instances of application, and those questions are therefore in need of similar responses in the $PM_{2.5}$ and the Ozone case. Accordingly, my responses to his questions on the $PM_{2.5}$ ISA should also be taken to apply here.

As was the case with the $PM_{2.5}$ questions, Dr. Cox raises an overarching fundamental concern about the sufficiency of how causality is understood and its properties characterized and applied, accompanied by a long set of very particular questions about whether these overarching concerns affect many particular details of the conclusions the document tries to justify. To address each of the particulars fully would require a quite massive dissertation on the general questions of how causality can be conceived of, recognized, and applied to support a regulatory analysis, and then a large further set of dissertation chapters – each one quite a task in itself – to apply this general discussion to its impact on all the particular specific inferences Dr. Cox asks about.

I fear that the time and effort available to me as a quick-response consultant does not allow such a deep dive. (It would take a lot of study for anyone not already focused on these issues as a main career path.) I am forced to try to address the thrust of the big questions to the extent I am able with my existing understanding and then briefly comment on selected application questions.

As I suggested in the comments on the $PM_{2.5}$ ISA, the issues raised are legitimate concerns, and what is needed is not quick responses from an array of commenters to a lot of particular instances of the larger questions, but rather a full taking on of a deeper look into causal inference as it can and should be applied in the evaluation of criteria air pollutants. This seems a topic for an appropriately chosen and charged NAS committee, with the time and expertise to bring epistemological principles, insights, and approaches into a general approach to assessment of the roles of particular pollutants in affecting public health, based on the kinds of data that are actually available.

Collectively, the questions raise the issue as to whether the interpretation of studies as evidence of causal influences of the agents under review with regard to potential health impacts have been rigorously pursued. Dr. Cox's questions raise the problem of whether appropriate distinctions are being drawn in the ISA analyses between observations of an agent's apparent association with toxicity outcomes of interest (which association one might *hypothesize* to be attributable, at least in part, to an actual causal role of the agent vis-à-vis the outcome) and a conclusion that the agent does indeed have such a causal role. There is the further issue of distinguishing the general ability to influence a toxicity outcome in some way under some circumstances (a qualitative property of the agent vis-à-vis its potential targets, potentially contingent on other factors, but asserted as the existence of potential for action divorced from particulars), and the inference that the agent has indeed acted causally in particular observed circumstances. The key question for the ISA application is really the reverse of this last one: given a set of particular observations about outcomes of interest and associated exposures to the agent, what kinds of associations – and what analytical approach to those associations – can justify an assertion that the agent operated as a causal influence in contributing to the bringing about of those outcomes? This needs to be distinguished from the further question as to whether the agent would be expected to

exert causal effects in *other* circumstances that will not share all the particulars of the observed case – that is, is the assertion of a causal effect seen in some studies can be generalized as a basis to predict that causal effects would be expected from other exposures? This in turn is distinct from the question of whether, if such qualitative prediction is deemed possible, the *magnitude* of the effect can be measured from observations to predict the magnitude of effect in hypothetical future exposure settings where other circumstances might differ.

My overarching reaction to these large questions is that Dr. Cox's concerns have merit. The treatment of "causation" in the ISA tends to be fuzzy – fuzzier in some places more than in others – about the distinctions noted above. If patterns of association are plausibly explained by underlying causation, this is taken as sufficient evidence for such causation (when one should actually be comparing the hypothesized causative actions against other competing explanations for the patterns), and if such causation is inferred, it is taken to be universal, applying to other settings on the usually poorly stated (much less justified) presumption that the causation is universal and largely independent of other circumstances. The measures of association are too easily taken as measures of the magnitude of cause, and these magnitudes applied to other settings without due consideration of how they may be contingent on differing circumstances.

This said, it should be clear that there is no ready "fix" to the issues raised. There certainly are procedures that can help identify and disentangle the various influences and contingencies -- and their use should be encouraged – but none will be absolutely definitive. There is no ideal analysis that can unambiguously characterize casual properties and their functional dependence on circumstances, such that an unambiguous and unquestioned prediction of the consequences of future exposures can be estimated without risk of tripping on one or several of the challenges that Dr. Cox's questions raise.

Clearly, one should attempt to identify and adjust for confounders. Directed graph analyses can help sort out questions as to which variables act directly *versus* through their effect on other variables. Experiences with outcomes that differ in their input patterns, along with controlled experiments, can help identify what factors can or cannot reasonably be expected to influence or interact with one another regarding outcomes. One should think through inferences to see where simplifying assumptions might be covering important influences.

The approach, then, should be to be aware of the pitfalls and challenges, and to beware of being naïve about the difficulty of sorting out the existence of some causal effects or about the challenges of asserting general properties from particular observations and asserting that they would apply, independent of other influences, in future settings that one wishes to predict to address the controls on exposure that would "protect public health with an adequate margin of safety." Just because a potential interpretation problem cannot be definitively eliminated, it does not mean that the information is meaningless or that there is no useful evidence about potential effects. But at the same time, the inability to eliminate a potential problem that cannot definitively be demonstrated does not mean that the problem does not exist or can simply be ignored.

It seems that much of the problem arises because of the tendency to think of a "cause" as both necessary and sufficient for the outcome. "Necessary" in that any effects seen are presumed to have come from the potential cause of interest (and not from other influences), and "sufficient" in that that cause alone is all that is needed, with no contingency on other factors as contributors or facilitators. This simple notion of "cause" will not really do for the kinds of effects pollutants can have on health.

Most pollutant effects on health will be INUS (Mackie 1974) – individually insufficient but (nonredundantly) necessary parts of a process that collectively is unnecessary but sufficient.

- That is, "insufficient" because other elements need to be present (behavior causing uptake, individual susceptibility is sufficient, etc.); but "necessary" (process will not happen without it, i.e., the "efficient cause" or "that without which"); "sufficient" because together the necessary element and its other components are enough to result in the effect; but "unnecessary" in that this collection of components is not the only way to get the effect (there are other causal pathways that can get to the same end result).
- Thus, the empirical association of a putative cause with effects is not dispositive for concluding it has a "causative" role, since (1) other, separate causative processes could have been responsible for the effects of interest (such that seeing the effect does not unambiguously identify the cause of it), and (2) the role of the putative cause – even if true – is contingent on other enabling features of the set of conditions (such that failing to see the effect after the putative cause could be explained by other needed conditions being missing).

What we usually intend to mean by something being "causative" is that it is an "efficient cause" in Aristotle's sense – the actor that imposes on other elements and, as a consequence of that instance of imposition, precipitates a change in their states. It is important to distinguish between (1) the assertion that this is an ability that an agent may exert (in some circumstances, i.e., a property of the agent's potential with respect to the states it might alter) and (2) the assertion that the agent has in fact led to a specific observed alteration of state (an assertion that, but for the presence of the agent, the particular alteration would not have occurred). These are both distinct from (3) the assertion that, whenever the agent encounters an appropriate target, it will act to effect its change – that is, the cause is by itself "sufficient" (not dependent or contingent on other circumstances; and (4) the assertion that an observed change can only be due to the specific cause (that is, that the cause is "necessary" for the effect, such that the observation of the effect ensures that the cause operated, with no other means available to bring about the effect).

Broadly speaking, there are two kinds of complications that arise: (1) the existence of potential confounders and varying influencers of outcomes, and (2) the further potential complexity arising from patterns of how various factors interact. The most obvious is the role of potential confounders. Several of the criteria air pollutants cause effects on the same endpoints, and the levels of these pollutants is highly correlated in space and time. Any analysis that does not adjust for such exposures will falsely attribute the entire effect to the one pollutant chosen as the "exposure" variable. In addition, non-pollutant factors can have a role, notably, the potential impact of differing weather (especially temperature) which can plausibly play an independent role, modify the impact of pollutant levels, and be correlated with those levels. To an extent, confounding can be mitigated by adjusting for the levels of the potential confounders in the analysis, but this can only go so far in eliminating the problem. One must recognize the factors in the first place, and many influential factors that ought to be adjusted for may be unrecognized or have insufficient data.

More than this, any adjustment itself makes untested assumptions about independence of the several influences and how they might interact to produce effects. Even if one has allowed for the correlated appearance of copollutants, to gauge their joint impact requires assumptions about whether they independently affect the outcome variable or somehow interact synergistically. In real populations, the values of the important variables will vary among individuals, and the particular combinations experienced by different people will differ, so what is observed is some kind of average over all the possibilities present in the measured population – averaging which can obscure the role of potentially important interactions. Any tractable analysis must make assumptions – usually untestable – about how

observed averages over person-to-person differences in levels, differences in combinations of exposures, differences in susceptibility, differences in the timecourse of exposure, and so on combine to yield an effect. If all the effects are independently acting and linear, then one can sum up the effects and express each as the average effect from average exposure (since for linear functions, the averaged output for variable inputs is equal to the function's value for the average input), but if they are not, it will depend on how they are not.

Clearly, these are not new concerns – they have been central to discussion of epidemiological methods through the history of that discipline. As I have argued, the key is to avoid being paralyzed by them but not to do this by ignoring them. One has to attempt to gauge the extent to which the conclusions from simplified but tractable analyses might be skewed by known or reasonably suspected complications, and then treat these insights as insights into uncertainty, to be factored in when one applies the analyses to decisions. Tools exist to help sort out the difficulties, and they should be used, but in the end, forthright expert judgment needs to be employed and (especially in the context of public-policy decisionmaking) the reasoning behind uncertainty characterizations being made explicit.

The way these issues affect the "causality" problem is that the complexities of context and potential interactions are what lead to the difficulty in inferring between the different aspects of causality – ability to affect, responsibility for observations, ability to generalize effect to other settings, ability to measure the impact of an effect, and ability to generalize that magnitude to other settings. They also affect the certainty with which causal determinations can be made, constituting the things that need to be thought through in characterizing uncertainty in any causation assertions.

As to the specific questions, I can offer the following.

Question 1 – It is not possible to make an undisputable, totally sufficient conclusion of interventional causation, because of the INUS nature of the possible causative processes. That is, there is no absolute necessity or sufficiency with which to show such causation by direct means (by giving or withholding the putative cause and seeing or not seeing the effect with absolute consistency). This is not to say that a hypothesis of interventional causation could not be proffered and examined against all the data that would be expected to show such an effect, paired with alternative explanations for how the effect pattern might arise from other plausible causes, and their comparative support compared. But it would have to acknowledge the possibility of other causes for the observed effects and of contingency of the interventional effect on other circumstances. This question of interventional causation, even if solved, does not release one from the challenges of sorting out what sets of conditions may contribute as well.

Question 2a – "a formal causal framework" implies that a path to certain determination is available, and I do not think one is. That said, one could establish a more explicit framework for assessing causation, and the NAS committee suggestion above would be a good way to pursue such general guidance – probably better than trying to write one in to an already existing ISA.

2b – No, as noted above, the naïve notion of causation sort of implies both necessity and sufficiency, but a more appropriate INUS view of how causes enter into health effect influence shows that the question is more complex.

2c – No, existence of a causal possibility does not guarantee that it would operate in any particular way. This is a further inference about its generalizability. The assertion that a cause operates independent of other varying potential influences is itself a strong conclusion that would be hard to defend in any absolute terms. So it is a question of working through how an effect seen in one setting may or may not

show up in others, with an assessment of any hypothesized independence-from-setting properties needed.

2d – Yes, it can be incorrect. It cannot be surely correct, so the question is how usefully to characterize the uncertainty and also how to act in the face of such uncertainty, given that failure to prove the correctness cannot be avoided, but it also does not prove that the causal process does not operate.

2e – I have argued that this could be improved, though it is hard to spell out a remedy in the context of these questions.

2f – there are methods to try to uncover contingencies on other factors, by measuring them and using directed graph analysis. But even these depend on the possible factors being noted and measured, and on other factor being absent, and on assumptions about independence of factors or lack of modification of risks by variable factors from case to case being true, and there will always be possibilities that cannot be addressed by available data and possible analyses. The world is always more complex than any feasible model, and the differences can affect conclusions, so one is always assuming that the needed simplifications do not miss something that is key, with plausibility arguments coming in where data are insufficient to give insight.

2g – I have argued that a more rigorous conception of causation is indeed needed. Bradford Hill considerations and other such approaches are at best guides for fuller thinking through, they are not tests to be passed.

2h – as argued above, no.

2i – as argued above, no.

2j – as argued above, no.

2k – unclear, partly because the meaning of each category could be tweaked. What is clear is that the categories do not seem well suited for making the distinctions between existence of some causal property, the degree of independence of that property from other circumstances, and the possibilities and limits of generalization.

2l – see 2k

2m – this is a matter of being clearer about hones standards.

Question 3 – In the time available, I am not able to chase down and document a basis for all these particulars, and they are more about reviewing the results than commenting on the technical method. My answers above will have to be used to address these questions.

Mackie, John Leslie (1974). The Cement of the Universe: A Study of Causation*. Clarendon Press.*