

Replies to Reviewer 1

Reviewer's comments are in standard font.

Responses and changes to the manuscript are in italics.

General Comments

Source attribution and source sensitivity techniques are valuable tools for air quality planners to understand air quality model results and to design effective emissions control strategies. The author presents an innovative approach that adds to the extensive existing literature on these methods. However, as described below, it is not clear in the manuscript how this method differs from existing sensitivity methods and source attribution approaches. It is also unclear how this method would be used in an air quality planning context, so I recommend revisions to the manuscript to more clearly explain this path-integral method (PIM) and to illustrate its use in the air quality planning context.

As described in the manuscript, a variety of approaches have been used to identify emissions source categories that are important contributors to ozone and other secondary pollutants. These methods can be broadly grouped as either model sensitivity methods or source attribution approaches. Sensitivity methods include forward sensitivities ("brute force" sensitivities and the Decoupled Direct Method (DDM)) and backward sensitivities (adjoint methods). Sensitivity approaches evaluate effects of changes in emissions on ozone or other pollutants relative to a base case model simulation. Source attribution methods rely on tracer species and/or evaluation of mass budgets and are used to evaluate sources that contribute to ozone or other pollutants in a particular model simulation.

Both sensitivity methods and source attribution methods have limitations, and neither fully addresses the needs of air quality planners who are tasked with identifying the most effective combination of emissions controls that demonstrate progress in reducing air pollutants (while avoiding possible dis-benefits of NO_x control) and that ultimately attain national ambient air quality standards. The key limitation of sensitivity methods is that ozone can have either positive or negative sensitivity to changes in emissions, and the magnitude and sign of the sensitivity depends on both the size of the emissions reduction and the sequential order in which different sources are controlled. Source attribution studies address this limitation of sensitivity methods by evaluating the contribution of each emissions source in a particular scenario, typically one that represents current conditions or an historical pollution episode. Thus, source apportionment methods can identify the largest contributors to ozone under current conditions, and this is useful for identifying and prioritizing sources to control. However, source apportionment methods do not predict the sensitivity of ozone to emissions controls in a future scenario because source attribution in the base case does not account for non-linear chemistry effects of emissions changes in the future case. Typically, air quality planners use source apportionment methods to identify potential emissions sources for control and then perform additional model sensitivity simulations to evaluate specific control scenarios

This manuscript describes a novel approach for using sensitivity simulations to assess how ozone and other pollutants respond to changes in emissions. The PIM method is designed to use a finite set of model sensitivity simulations to systematically represent the range of model response to control of each emissions category ranging from 0 to 100% control, while also representing the variation in model response depending on the order in which source categories are controlled. The paper illustrates the PIM method using a highly simplified box model scenario with two layers. The author notes that a

limitation of the method is that a large number of model sensitivity simulations are required and that this would result in large computational cost. Another limitation not identified in the paper is that presentation of the PIM results could also be challenging for a realistic model scenario. It would be more useful to illustrate the application of this method using a realistic model simulation. My initial reaction is that this method would be challenging to use in an air quality planning context, and that it does not provide insights that are not already available from conventional sensitivity and source apportionment methods. Therefore, I recommend that the manuscript be revised to illustrate the PIM method using a more realistic model scenario using a 3-dimensional photochemical model for an historical ozone episode used in an air quality planning context.

Response: The PIM is a new method that introduces new concepts and is consequently more difficult to explain and understand than a modification of an existing method. This paper is intended to explain the method, indicate good options for integrating the sensitivities numerically, and illustrate the method for a chemically realistic model. The model used here contains the full CB6 chemical mechanism used in the 3-D model and therefore all the nonlinearities in the 3-D model. Thus, the work provides a valid test of whether the PIM can accurately account for the effects of the nonlinear chemistry. The PIM has also been used with the full, 3-D version of CAMx in an application to the eastern U.S. This application is reported in another paper that has been submitted to a journal. Combining the two papers would result in a long paper that would likely be less, not more easily understood.

Manuscript change: A reference to the paper describing use of the PIM in a 3-D model has been added at the end of the Introduction and at the end of the Conclusions.

Response: In the 3-D application, the PIM required between 2.5 and 3.0 times the effort of the brute-force (zero-out) method. This is larger than the effort for some other methods, but not so large as to be unmanageable.

Response: If only one emission-control scenario is studied, presenting the PIM results is no more challenging than presenting the results from other source apportionment methods, namely providing the source contributions for specific receptor locations or geographic areas. If one chooses to evaluate multiple emission control strategies, then there will be a different set of source contributions for each control strategy of interest. This adds complexity but also provides new, useful information on how the importance of sources varies depending on the control strategy.

Response: The PIM does provide insights not available from other methods. It provides contributions for anthropogenic sources that sum up to the difference (Δc_i) between a simulation with and a simulation without the anthropogenic sources. This is not a requirement for other methods and is generally not true for results from other methods. If the sum of the anthropogenic contributions is greater than (less than) Δc_i , the results indicate that reducing anthropogenic emissions will produce greater (lesser) pollutant reductions than will actually occur (based on the model predictions). This can obviously be misleading for air quality planning purposes. The PIM also provides source contributions that effectively are averaged over a range of emissions and chemical regimes. Other source apportionment methods and sensitivity methods only use the emissions and chemistry in the base case, without any emission reductions. Lastly, the PIM provides source contributions for all species in the chemical mechanism simultaneously, e.g., O_3 , NO_2 , and air toxics, without additional assumptions exterior to the model, e.g., when the chemistry is VOC-limited vs. NO_x -limited. This is not true for source apportionment methods based on tracers.

Manuscript change: The following paragraph has been added to Section 2.1:

“Because the sensitivities are integrated over the path P in Eq. (2), the PIM considers a range of chemical conditions in calculating the source contributions, from zero to the full anthropogenic emissions in the base case. Methods based on tracers or a Taylor series expansion (e.g., with first- and second-order sensitivities) use only the emissions and the chemical conditions of the base case. Thus, the PIM provides source contributions that are averaged over the emission-control scenario, not specific to the base case.”

It would also be helpful to compare PIM with the high-order DDM method which accounts for some nonlinearity in the photochemical reactions. Can the high-order DDM provide the same information that PIM calculates using multiple sensitivity simulations?

Response: The high-order DDM (HDDM) does not provide the same information as the PIM. The HDDM involves constructing a Taylor series expansion through 2nd order for the emissions (and the chemical regime) in the base case, and the expansion is then used to extrapolate to lower emissions. However, the extrapolation is generally accurate for no more than a 50% reduction in anthropogenic emissions. At larger reductions in anthropogenic NO_x emissions, the chemical regime changes to a high VOC/NO_x ratio due to the large biogenic VOC emissions, the chemistry becomes progressively more sensitive to NO_x emissions, and the sensitivities calculated by HDDM for the base-case emissions are no longer accurate. Therefore, the sum of the source apportionments from the HDDM does not agree well with the anthropogenic increment of the O₃ concentration. The PIM avoids this problem by using sensitivities over the full range of anthropogenic emissions (0% to 100%). Furthermore, the HDDM approach calculates 2nd order cross sensitivities between different sources, but, as discussed in Section 2.2, the cross terms are not assigned to individual sources in the source apportionment. The PIM effectively assigns these cross terms to sources based on the emission-control strategy. Section 2.2 and the Supplementary Information show how this could be done if a Taylor series expansion of 1st order sensitivities is accurate (though, as indicated above, such an expansion actually has a limited range of applicability).

As a general comment, the description of the method seems to be overly abstract and it would be helpful to explain in simple language the physical significance of terms such as the path variable and hypercube. Also see comments on page 5 below.

Manuscript changes:

The sentences in Section 2.1 where u is introduced are revised to: “The path P can be described via a path variable u that describes position along the path. Each λ_m is a function of u , such that as u varies from 0 to 1, each $\lambda_m(u)$ also varies from 0 to 1, and the path P defining the changes in anthropogenic emissions is traced from the background case to the base case in the M -dimensional space of the scaling parameters λ_m . “

The second sentence in Section 2.2 is revised and a new sentence added: “This is a path along the edges of a hypercube in Λ -space. (The hypercube defines all possible emission-control strategies, contains M axes, one axis for each λ_m , and includes all values of λ_m from 0 to 1.)”

Specific Comments

Page 3, lines 10-17: The text in this paragraph is difficult to follow, see comments below: “If the anthropogenic increment is allocated to sources, the PIM requires that the base-case concentration minus the sum of the anthropogenic source contributions equals the background concentration. Other methods do not have this requirement, and thus may over- or under-allocate the anthropogenic

increment to the anthropogenic sources and also allocate a concentration to the background sources that is not the concentration from a simulation with only background sources included.”

Are there cases in PIM in which the anthropogenic increment is not allocated to sources? How does the PIM treat apportionment if the anthropogenic increment is negative? Does this method assume that increments are always positive? Also, it is not necessarily correct to state that other approaches “over- or under-allocate”. Given the constraint of accounting for negative sensitivity of O₃ to precursors and accounting for the contribution of both VOC and NO_x to O₃, each method adopts a unique strategy for mass attribution. Allocation can be internally consistent with the adopted strategy and therefore technically accurate, and yet provide estimates that differ from other apportionment approaches. Instead, the author might argue that the strategies adopted in other apportionment approaches are poorly understood and result in incorrect interpretation of the results, or are incorrectly implemented and therefore produce inaccurate results. However, more description and analysis of results from other apportionment methods is needed to support such a conclusion.

Response: As discussed in the 3rd paragraph of Section 2.1, the PIM can also allocate the total concentration, not just the anthropogenic increment. However, to allocate the total concentration, the initial concentrations, boundary concentrations and emissions from all sources (biogenic and anthropogenic) must be included in the analysis.

Response: If the anthropogenic increment is negative, then some anthropogenic source contributions may be positive, but the sum of all the anthropogenic contributions will be negative (Eq. (1)).

Response: The method does not assume that anthropogenic increments are always positive. The O₃ increment is negative from hours 0 to 12 in the simulation (Figure 2) and the NO_x source contributions are negative during these hours (Figure 4).

Response: To my knowledge, none of the other methods requires that the base-case concentration minus the sum of the anthropogenic source contributions (call this difference δ) equals the background concentration. If δ is less than the background concentration, then the method assigns too much importance to the anthropogenic sources and will give the impression that reducing anthropogenic emissions will lower the pollutant concentration more than will actually occur (over-allocation of the anthropogenic increment to the anthropogenic sources). Similarly, if δ is greater than the background concentration, the method is assigning too little importance to the anthropogenic sources (under-allocation of the anthropogenic increment). I think that the text is a fair statement of this important difference between the PIM and other methods, given the definition of over- and under-allocation used here. However, to explain the issue more clearly, changes have been made to the Introduction and Conclusions.

Manuscript changes:

Section 1. Introduction, p. 3, lines 10-17 are revised to: “An important advantage of the PIM is its ability to allocate to sources a concentration increment, i.e., the difference between two simulations (base and background cases). If the anthropogenic increment is allocated to sources, the PIM requires that the base-case concentration minus the sum of the anthropogenic source contributions equals the background concentration. Other methods do not have this requirement, and thus may ascribe too much or too little importance to the anthropogenic sources.”

Section 5. Conclusions, first paragraph is revised to: “As shown in Sect. 4, the PIM can allocate the difference in concentration between two simulations to emission sources. Consequently, the PIM requires that the base-case concentration minus the sum of the anthropogenic source contributions (difference δ) equals the background concentration (within the accuracy of the numerical integration). Other methods do not have this constraint. If δ is less than the background concentration, then the method assigns too much importance to the anthropogenic sources and will give the impression that reducing anthropogenic emissions will reduce the pollutant concentration more than will actually occur (over-allocation of the anthropogenic increment to the anthropogenic sources). Similarly, if δ is greater than the background concentration, the method assigns too little importance to the anthropogenic sources (under-allocation of the anthropogenic increment). The PIM ensures that the anthropogenic increments to O_3 and the other species are neither over- nor under-allocated to the anthropogenic sources.”

Page 3: Equation 1 includes only first-order sensitivities of c_i with respect to the scaling parameters. A term is also needed to represent higher order sensitivities.

Response: Eq. (1) is exact as written. No higher-order sensitivities are involved because the integrals are over the variables (λ_m) with respect to which the derivatives are taken. Eq. (1) is a generalization to multiple dimensions of the familiar relationship: $f(b) - f(a) = \int_a^b \frac{df}{dx} dx$. In one dimension, there is only one path from a to b . In multiple dimensions, there is an infinite number of possible paths from the starting to ending point of the integration

Manuscript change: The first paragraph of Section 2.1 is revised to: “The PIM is based on an exact mathematical equation that is in itself not new. In particular, the equation is routinely used in thermodynamics (Sect. 2.3). However, the application of the equation to atmospheric modeling is new. The equation is the generalization to multiple variables of a familiar relationship for a single variable, namely that the integral of the derivative of a function ($\int_a^b \frac{df}{dx} dx$) is equal to the difference in the value of the function at the ends of the integration interval ($f(b) - f(a)$).”

How does this result differ from higher order DDM?

Response: See the comments above about HDDM.

Page 4, lines 20-22: “However, if all the source contributions and Δc_i are calculated, then Eq. (1) can be used to check the accuracy of the integration procedure. The integration procedure can be modified then, if necessary, so that the sum of the source contributions equals Δc_i within the desired error tolerance.”

Because models are not strictly mass conservative and are subject to numerical error, an approach is needed to avoid accumulation of error. Thus, a method is needed to prevent accumulation of error in the case where all source contributions are not calculated.

Response: The only difference between calculating all source contributions S_{im} and calculating a subset of them is simply that, for the omitted S_{im} , the sensitivities aren't calculated and the integration in Eq. (2) is not done. There are no changes to the sensitivities that are calculated, and there are no additional errors in the numerical integration procedure for the S_{im} that are calculated compared to the case where all S_{im} are calculated. (The DDM gives the same value for a particular sensitivity coefficient regardless of

whether only that sensitivity or multiple sensitivities are calculated.) Therefore, there are no new errors and no accumulation of error in the case where only a subset of the source contributions are calculated.

Page 5, lines 3-7: It seems very problematic that the source apportionment result depends on the order in which emissions sources sensitivities are calculated. Also, the definitions of path (P), path variable (u) and normalized distance (s) are not clear. What is the physical significance of normalized distance or absolute distance along P?

Response: The source apportionments depend on the path P but not on the order in which emissions sensitivities are calculated. The apportionments depend on the path because the chemical environment changes depending on how emissions are controlled. The simplest example is the special case of successive zero-out of sources, Section 2.2. The apportionments depend on which source is removed first, which second, etc.

Response: As indicated above, the sentences in Section 2.1 where u is introduced are revised to clarify the definition of u and P.

Response: There is no special physical significance to the normalized distance s. Use of s rather than u is just a change of variable that can make it easier to understand where the Gauss-Legendre integration points will be, ensure a more uniform distribution of the points, and thus may make the integration more accurate.

Manuscript change: The sentence prior to Eq. (4) is revised to: "Changing the integration variable from u to s, the source contribution becomes"

One of the key limitations of source sensitivity methods is that sensitivities are not additive, and that that O3 can have negative sensitivity to precursors in some cases. A key motivation for source apportionment methods is to estimate the actual mass contribution of a source to O3 rather than the sensitivity of O3 to that source. It would seem that the PIM methods suffers from the limitation of sensitivity approaches and does not provide a mass attribution estimate that is unaffected by nonlinear sensitivities.

Response: The PIM provides an allocation of pollutant concentrations to emissions by integrating first-order sensitivities over emission levels. Because it integrates over the emission levels, it includes all nonlinear effects, and the integral of each sensitivity is a mass amount, not a sensitivity. The source contributions are additive, as shown by Eq. (1).

Page 5, lines 23-24: "The simplest and shortest integration path, termed the diagonal path, is defined by $\lambda_m = u$, all m. This is a straight line from $\Lambda = 0$ to $\Lambda = 1$ along which the emissions from all sources are reduced or grown by the common factor u."

Previously "u" was defined to be the path variable, which was unclear. If u is a factor by which emissions are adjusted, this could be explained more clearly. It is confusing that an emissions change factor is described as a distance.

Response: As indicated above, the sentences in Section 2.1 where u is introduced are revised to clarify the definition of u. The λ_m scale the emissions, u describes position along the emission-control path, λ_m is a function of u, and for the diagonal path $\lambda_m(u)$ equals u. For other paths, $\lambda_m(u)$ does not equal u, e.g., Eqs. (6, 7).

Manuscript change: The two sentences are revised to: “The simplest and shortest integration path, termed the diagonal path, is defined by $\lambda_m(u) = u$, all m . This is a straight line from $\Lambda = 0$ to $\Lambda = 1$ along which the emissions from all sources are reduced or grown by the common factor $\lambda_m(u) = u$.”

Page 8. “Analogy in thermodynamics”. Suggest deleting this section as it does not seem relevant and is not helpful for illustrating the PIM method.

Response: Reviewer 2 asked for examples of other problem domains where Eq. (1) has been used, so it seems important to keep this example.

Page 14, lines 7-10: “The PIM allows source contributions to be either positive or negative. If the secondary pollutant formation is inhibited by emissions of some species, source, or geographic area, the sensitivity to these emissions will be negative for at least some values of the scaling parameter m , and the integral in Eq. (2) may be negative.”

The above statement highlights the difference between sensitivity and source apportionment methods. Negative “contribution” indicates that this is a sensitivity method, not a mass attribution method. NO_x emissions can contribute to ozone production even when ozone has a negative sensitivity to changes in NO_x emissions, and source attribution methods such as OSAT in CAMx are designed to quantify the mass contribution of NO_x to O₃. Thus, the PIM method is not quantifying the mass contribution to ozone production (in the sense that it evaluated in a source apportionment approach), rather, it is characterizing the negative sensitivity of ozone to NO_x.

Response: The anthropogenic increment can be negative. (See O₃ in Figure 2, hours 0 to 12). This means that removing the anthropogenic emissions actually increases the O₃ concentration, which is due to the titration of O₃ by NO emissions. If the anthropogenic source contributions are all positive, then the implication is that eliminating the anthropogenic sources will reduce O₃, which is incorrect. The PIM integrates the sensitivities over the anthropogenic emissions and accumulates the effects of both positive and negative sensitivities to the emissions. The integral of the sensitivity is no longer a sensitivity, but a source contribution. If the anthropogenic increment is negative, then the sum of the anthropogenic source contributions will be negative, which is a consistent and correct result.

Page 14, lines 23-24: “The concentrations in the background simulation can be determined by an actual simulation or by subtracting the source contributions from the base-case concentrations.”

The solution for source contributions is non-unique, i.e., the solution depends on assumptions made in the order in which sources are evaluate. Therefore, subtracting source contributions from the base case does not provide a unique estimate of background concentrations. The only reliable modeling approach to estimate background concentrations is to perform a model simulation that does not include anthropogenic emissions

Response: The source contributions are not unique, but the sum of the source contributions is always the difference Δc_i in Eq. (1).

Manuscript change: The sentence is revised to: “The concentration in the background simulation can be determined by an actual simulation or by subtracting the sum of all the source contributions from the base-case concentration.”