

Dear Editor,

We thank you for the opportunity to revise the manuscript and to address the Reviewers' comments. In the revised version, we propose the improved title "*Causal inference for quantifying chemical–dynamical pathways controlling tropical middle stratospheric ozone variability*" to better reflect the focus of the manuscript.

While the overall methodology remains unchanged, the manuscript now focuses on a dominant chemical–dynamical pathway that influences tropical middle stratospheric ozone (O_3) over the longer 2004-2021 period, rather than analyzing shorter subperiods. In the revised manuscript, the main objective is to quantify the mechanisms controlling O_3 variability on monthly timescales, which can in turn modulate trends over limited time periods. Therefore, the Introduction motivates this study by discussing ozone changes and trends reported over shorter time periods. This is now clarified in the revised manuscript. We also performed an additional causal analysis to address the concerns raised by the Reviewers, which we address in detail in our responses.

Best regards,
Evgenia Galytska
On behalf of the co-authors

Review of Galytska et al 2025: “Causal inference for stratospheric chemistry: insights into tropical middle stratospheric ozone variability” submitted to ACP

December 3, 2025

We thank the Reviewer for providing the feedback on this manuscript for the second time. We revised the motivation and, consequently, the focus of this manuscript, resulting in updated results. While the overall methodology remains unchanged, we extended the discussion of important aspects, such as the necessity for detrending timeseries (see Sect. 3.1), interpretation and discussion of the N₂O-NO₂ causal relationship (see Introduction, Sect. 4.1). We would also like to highlight that we now propose an improved title, “*Causal inference for quantifying chemical–dynamical pathways controlling tropical middle stratospheric ozone variability*”, to better reflect the focus of the manuscript. We addressed the comments below one by one in blue. We use the notation **P L** to refer to the changes in the revised manuscript on a specific **Page** and **Line**.

The submitted paper uses methods from causal inference to examine short-term trends in ozone in the tropical middle stratosphere since 2004, with a particular focus on the difference between trends during the periods 2004-2011 and 2012-2018. Sophisticated methods are used that bring together causal inference, observations, prior physical understanding, and chemical transport models. The paper has a direct message about the importance of transport-induced variability in N₂O and NO₂ and a ‘meta-message’ promoting the usage of causal inference methods in stratospheric photochemistry.

Note that this is the second time I have reviewed a version of this paper, having reviewed an earlier version submitted to a different journal. The paper has been revised since my first review, and I have had the opportunity to read it with fresh eyes. Although two major comments (and all of my minor comments) were addressed since my first review, others remain to be addressed. I also raise some new concerns upon this fresh reading.

As noted in my first review, broadly, this paper does a commendable job bringing together a range of methods to explore trends in mid-stratospheric ozone. A strength of the paper is that it clearly explains the causal inference mechanisms that it is deploying, including with the help of an illustrative toy model. These methods are used to support the key result of the paper, which is essentially a simple mechanism previously explored by some of the same authors (Galytska et al., 2019): in the tropical mid-stratosphere, where the dominant sink of ozone is catalytic destruction by NO_x, stronger residual

upwelling has led to less NO_x and more O_3 (and weaker residual upwelling leads to more NO_x and less O_3). This could explain some decadal variability in ozone at these altitudes, which is well worth explaining.

The paper supports this simple mechanism with sophisticated methods, which it illustrates as a proof of concept, and with some auxiliary nontrivial results. Although I am confident in the value of explaining ozone variability at these altitudes, and this mechanism seems quite plausible, I still have significant concerns with specific methodological choices made in this paper, which presently do not allow it to add as much value to current understanding as it should. One of my major concerns from the first review has been fully addressed, but the other three remain, and I must also introduce a fourth:

- A primary motivation of the study is a puzzling reduction of ozone from 2004-2011, yet this declining trend was not statistically significant, and it is excluded from the analysis by detrending in the data preprocessing step.

This has been fully revised. We no longer analyze two subperiods as the primary motivation of the manuscript and instead focus on the long 2004–2021 period. The Introduction has been revised to address the importance of understanding the mechanisms governing interannual ozone variability in the tropical middle stratosphere.

- The directed acyclic graph implies that low amounts of N_2O “cause” high amounts of NO_2 (a Simpson’s paradox whereby an increase in reactant is associated with a reduction of reaction product), yet this seems ill-suited to capture the response to variability in the source of N_2O .

We understand the Reviewer’s concern. The negative contemporaneous link from N_2O to NO_2 is a result of their opposing response to transport-driven variability, rather than a direct chemical effect. We now make sure that this is properly discussed in the manuscript. For example, in the Introduction we start with the discussion of the negative N_2O - NO_2 correlation on **P3 L72-74**: “The strong anti-correlations between N_2O and NO_2 (panel b) result from their opposing response to transport-driven variability, where, e.g. enhanced upwelling increases N_2O while reducing its chemical loss, which is crucial for NO_2 production (see further discussion in Sect. 4.1).”

Section 4.1 addresses this link while discussing Fig. 3a on **P10 L239-246** as follows: “The relationship from N_2O to NO_x (labelled B) is negative, despite N_2O being a source of NO_x . This apparent contradiction is an example of Simpson’s paradox (Blyth, 1972) and arises because tropical residual velocity w^* acts as a confounding dynamical process, leading to an anti-correlation between N_2O and NO_2 . Namely, slower (faster) upwelling results in lower (higher) N_2O concentrations and consequently longer (shorter) N_2O residence time in this

region, which allows more (less) time for the photochemical production of NO_x from N_2O . Consequently, higher (lower) NO_x levels lead to lower (higher) O_3 concentrations via the NO_x -catalyzed ozone destruction cycle, resulting in a negative relationship (labelled C, see Crutzen, 1970). Table 1 summarizes the discussed chemical-dynamical relationships in the tropical middle stratosphere as depicted in Fig. 3a.”

Also, in the discussion of Table 1 (**P11**), we now explicitly address this point as follows: “ N_2O is the primary source of NO_2 via the slow reaction $\text{N}_2\text{O} + \text{O}(^1\text{D}) \rightarrow 2\text{NO}$, followed by the rapid reaction $\text{NO} + \text{O}_3 \rightarrow \text{NO}_2 + \text{O}_2$. As N_2O increases due to faster circulation, its residence time decreases, and therefore NO_2 concentrations decline (there is less time to produce NO_2). The resulting link is therefore negative. In the absence of dynamical variability, the relationship would appear positive, reflecting only the underlying chemical production of NO_2 from N_2O ”.

In addition, the sensitivity test applied to the observations reveals an interesting feature related to the N_2O – NO_2 connection, which we further discuss below and in the manuscript on **P14 L319-323** and Fig. 5.

- The paper emphasizes that small (0-1 month) differences in the lags between variables constrains the upwelling, but it is not clear why.

Since we do not analyze separate subperiods anymore, this small lag difference is no longer relevant and has been removed from the revised manuscript.

- Is it appropriate for this method to omit a temperature-mediated pathway between upwelling and ozone? (Stronger upwelling can lead to cooling, which can increase ozone)

We have now estimated the connections including temperature in the causal graph. We performed additional causal discovery experiments to estimate a temperature-mediated pathway between upwelling and ozone, but unfortunately, we did not obtain a robust and statistically stable direct causal effect for this pathway. One of the possible reasons could be the diagnostic artifact, because w^* is a derived Transformed Eulerian Mean (TEM) variable that depends on thermodynamic fields (including temperature) and is not an independently observed quantity. We further discuss it below (see response to point 1.4) and in the revised manuscript in Sect. 4.2 **P13 L288-294** as follows: “However, temperature-mediated effects could amplify the apparent strength of the connection from NO_2 to O_3 , since enhanced upwelling induces both cooling (increasing O_3) and reduced NO_x species. To assess this, we performed additional analysis, including temperature anomalies in the tropical middle stratosphere. The inferred graph structure was not robust, likely because w^* is a derived diagnostic that depends on thermodynamic fields. Removing w^* altered the parent structure and prevented a direct comparison of direct causal effects

shown in Fig. 4a. We therefore interpret the identified NO₂-mediated pathway as the dominant mechanism, while acknowledging that temperature-related effects may potentially project onto this link”.

Because these concerns relate to fundamental methodological and interpretive choices in the manuscript, I will be recommending that this paper is Rejected with Encouragement to Resubmit. I see great promise in these methods and in the style of causal inference matched with observations and comprehensive models in this work, so I believe that, conditional on the resolution of these concerns, this paper has a bright future.

1 General comments

1.1 Challenges understanding the motivation and its relation to methodological choices

I see challenges in understanding the intended scope of the paper. What physical drivers of variability are intended to be included versus excluded? How general are the results intended to be, and how general are they? The concerns raised in this section are the most important, as I believe the paper needs to improve the alignment between its stated goals, its methods, and its conclusions.

To begin, I am concerned by how the paper treats ozone trends. It divides the period 2004-2018 into two subperiods (2004-2011 and 2012-2018), and it is stated (e.g., in the abstract at Lines 5-6) that ozone decreased during the first subperiod and increased during the second subperiod. But, Figure 1a appears to instead show a very noisy time series with no obvious trend within each subperiod. Any significant trend that was found would likely suffer from endpoint problems. The variability is dominated by the QBO, and during an 8-year period, the start and end phase of the QBO could be very important for trends.

We refrain from analyzing the subperiod as the major focus of the manuscript and instead focus on a single long 2004–2021 period. We therefore removed from the manuscript Fig. 1a, which previously showed yearly ozone anomalies for the two subperiods.

It is further confusing how this questionable trend enters into the arguments of the paper. The Introduction is mostly dedicated to discussing decadal ozone trends, which gives the (perhaps mistaken) impression that the paper will be about processes that drive decadal ozone trends. However, in the data preprocessing step, the subperiod time series are detrended before performing causal inference (Lines 144-145). Thus, the method removes the 8-year trend that motivated the consideration of this time period,

and focuses instead on residuals from that trend.

We have revised the Introduction, made it more concise and reduced the discussion of tropical ozone trends. However, the remaining discussion on tropical *middle* stratospheric ozone trends serves as the motivation to quantify the mechanism that governs interannual ozone variability. On **P2 L49-50** we state: “*While previously discussed O_3 trends motivate this study, our objective is to quantify the mechanisms that control O_3 variability on monthly timescales and thus, that can modulate trends over limited time periods*”.

There is no guarantee that residuals from a long-term trend will exhibit the same relationship between the variables of interest as does the long-term trend. Specifically, these residuals are probably dominated by the QBO, as the paper analyzes later, and it is not guaranteed (nor likely) that the relationship between upwelling and ozone (or mediating variables) is identical in QBO-induced variability compared to other sources of variability that might be implicated in long-term trends, such as rising CO_2 , rising N_2O , or declining ozone-depleting substances. This review will later return to some of these themes.

We now address this issue in the Introduction on **P3-4 L86-89** as follows: “*It is important to highlight that to ensure statistical consistency, this study analyzes detrended anomalies rather than long-term trends. Therefore, the results describe variability-driven processes and should not be interpreted as a direct explanation of decadal O_3 trends caused by externally forced long-term changes in CO_2 , N_2O emissions, or ODSs, but rather in the context of dynamical processes that are strongly influenced by the QBO*”.

Indeed, in Section 4.4, the algorithm is tested on data that has been stratified by QBO phase, so that the variability is not necessarily driven by the QBO anymore. Despite this section discussing extensively how this is a robustness test for the methodology (Lines 353- 357), my takeaway is that the DAG doesn't look very robust. It looks qualitatively different when QBO-driven variability is conditioned out (Figure 6). Comparing across the rows of Figure 6, it appears that the relationship between upwelling and N_2O vanishes in both subperiods at a given QBO phase. Thus, it would seem to me that the (unconditioned) relationship between upwelling and N_2O is dominated by QBO-driven variability. In 2012- 2018, the DAG has no linkages in common during easterly wind shear as during westerly wind shear, and neither of these reproduces the full unconditional fit.

We agree with the reviewer that the previously regime-oriented analysis did not look very robust and suggested that some vital links vanish under specific conditions/regimes. One of the reasons is the very narrow shear layer that we selected to

identify QBO regimes. We now calculate shear between 10 and 30 hPa for the process-oriented analysis. We also refrain from applying the causal discovery algorithm for individual QBO regimes, as a limited sample size reduces the robustness of the causal discovery algorithm. We, therefore, now focus on assessing how the previously detected connections vary with the QBO phase. This is now addressed in **P15 L335-337** as follows: *“To assess how QBO phase affects the strength of the identified relationships, we keep the causal graph inferred from the full period and re-estimate the link strengths separately for easterly and westerly QBO conditions. This approach isolates regime-oriented changes in coupling strength without altering the underlying network structure”*.

So, it seems that the QBO might be dominating the relationships between upwelling/N₂O/NO_x/ozone documented in this paper, and that these relationships might be in some sense particular to the QBO, rather than general to other mechanisms. Noting these limitations, I hope that the paper can be clarified on whether it is providing a general understanding of upwelling/N₂O/NO_x/ozone relationships applicable to trends vs. a specific understanding of upwelling/N₂O/NO_x/ozone relationships driven by the QBO. The scope of information provided by the methods should be aligned well with the stated goals of the paper, especially in the Introduction.

With the revised approach, Fig. 6 now shows how the strength of the specific links varies across different regimes. It further demonstrates that the chemical-dynamical mechanism driving tropical middle stratospheric O₃ variability remains robust across different QBO phases. The major findings of this analysis are summarised on **P16 L353-355** as follows: *“Overall, the sign of all connections is robust across datasets and QBO phases. In the observations the chemical links (from N₂O to NO₂ and from NO₂ to O₃) tend to strengthen during the westerly QBO phase.”*

1.2 N₂O source variability?

During my review of the prior submission, I raised concerns about structure of the directed acyclic graph (DAG) and whether low amounts of N₂O can be understood to “cause” high amounts of NO₂. Aspects of the discussion about this question have been revised, which I appreciate, including a reference to the idea that the negative correlation between N₂O and NO₂ is an example of a Simpson’s paradox (Line 242). I still have core concerns about how to understand the linkage between N₂O and NO_x in the DAG.

It seems that the linkage from N₂O to NO_x results from the superposition of (at least) two physical processes: (1) more source N₂O tends to causally lead to more NO_x (just as generally more of a reactant leads to more reaction product), and (2) more source N₂O tends to be statistically associated with less NO_x due to their mutual response to upwelling (stronger upwelling increases N₂O while decreasing NO_x). The

fact that these methods find a negative relationship between N_2O and NO_x must be that upwelling-mediated co- variations in N_2O and NO_x have dominated over variability in the source of N_2O over the historical period of interest. But even if the upwelling is dominant, might there be some role for source variability in N_2O that affects the DAG structure or coefficients? How would this be represented in your method? Would it tend to make the linkage between N_2O and NO_x less negative? These questions relate again to the first question about the generality of the results: is the structure and coefficients of the DAG reflecting a mechanism-agnostic relationship between N_2O and NO_x or rather a QBO-specific relationship? (I suspect the latter.) These considerations are further relevant considering the stated motivation of the paper, because in the Introduction, the paper refers to possible trends in NO_x from long- term changes in N_2O , but it is not clear that this method would be able to capture those, since the method isolates the negative relationship between N_2O and NO_x associated with QBO variability rather than the positive relationship that would be associated with an increasing trend in tropospheric N_2O .

We thank the Reviewer for this thoughtful point. While we briefly noted above that the negative contemporaneous link from N_2O to NO_2 reflects their opposing response to transport-driven variability rather than a direct chemical effect, we would like to briefly elaborate on this below.

We agree that two processes can influence the N_2O – NO_x relationship:

- (i) source-driven pathway and
- (ii) transport-driven pathway.

In the case of **(i) source-driven pathway**, if surface N_2O increases, more N_2O enters the tropical middle stratosphere and that can lead to more ($NO_y/NO_x/NO_2$) over long time scales, which leads to a positive relationship.

In case of **(ii) transport-driven pathways**, the changes in upwelling strongly affect both species, typically increasing N_2O (associated with younger air) while decreasing NO_x (via transport/chemistry), which produces an apparent negative relationship.

Our analysis is based on monthly anomalies in the tropical middle stratosphere, where variability is dominated by stratospheric transport changes (and QBO-related circulation), not by year-to-year changes in the surface source of N_2O . For this reason, the method finds a negative N_2O – NO_2 link, which we interpret as mainly reflecting the transport-driven pathway.

We completely agree with the Reviewer that if source-driven variability in N_2O were stronger, it could weaken the negative N_2O – NO_2 link (make it less negative or even positive) because it would add variability where N_2O and NO_x increase together. In our framework, this would be represented by an additional driver of N_2O that is independent of transport (for example, a slowly varying “ N_2O source” term). Since we do not include such a source term explicitly and we focus on detrended anomalies, our results should be understood as describing the dominant transport-driven variability, not the response to changes in surface emissions. To highlight that our study now

analyzes the variability-dominant relationship, the Introduction on **P3-4 L86-89** states the following: *“It is important to highlight that to ensure statistical consistency, this study analyzes detrended anomalies rather than long-term trends. Therefore, the results describe variability-driven processes and should not be interpreted as a direct explanation of decadal O₃ trends caused by long-term changes in CO₂, N₂O emissions, or ODSs, but rather in the context of dynamical processes that are strongly influenced by the QBO.”*

In addition to the changes made to clarify this point (see the beginning of this revision), we note that, based on Fig. 5, the sensitivity tests applied to the observations indicate that at a more relaxed significance threshold, the causal discovery algorithm detects a weak positive one-month lagged connection from N₂O to NO₂. This is physically consistent, as N₂O is a source of NO₂. However, this link is sensitive to the choice of significance threshold and is not statistically robust. We have included the discussion of this connection on **P14 L319-323** as follows: *“A further analysis of the sensitivity experiments reveals an additional feature in the N₂O-NO₂ relationship. A positive one-month lagged link from N₂O to NO₂ is detected across all tested τ_{max} but primarily at relaxed significance thresholds ($\alpha_{pc}=0.2$). For $\tau_{max}=1$, the link is also identified at $\alpha_{pc}=0.1$ and 0.05 . Such a positive lagged connection is physically plausible, as NO₂ is produced from N₂O, as discussed in Table 1, and a delayed response may emerge at the monthly timeseries. However, given its sensitivity to the choice of α_{pc} , this link cannot be considered a robust pathway”*.

1.3 Interpreting 0 vs. 1 month lags?

The paper ascribes significance to a seemingly small difference in the value of the lag between N₂O and NO₂ (0 months in 2004-2011 versus 1 month in 2012-2018). I am not sure whether this difference is significant or how it implies that upwelling has increased. I appreciate that a sensitivity test has been added in Appendix B, which shows that the 0 vs. 1 month lag difference is robust to some amount of variation in the hyper-parameters. Robustness notwithstanding, I am also concerned about the claim that a change in the lag between N₂O and NO₂ implies that upwelling has increased. First, upwelling is an input to the algorithm by which the linkages are calculated, so I am not sure whether the results of the algorithm can provide an independent constraint on upwelling. Please comment on this. Second, I am confused about whether the target to explain here is the average upwelling during the subperiod (implied in Lines 286-288) or the subperiod trend in upwelling (implied by the comparison with the LOTUS subperiod trends).

We now recognize that the previous formulation was misleading, since tropical middle stratospheric trends served as a motivation for the study. We have now fixed it in the Introduction **P2 L49-50** as follows: *“While previously discussed O₃ trends motivate this study, our objective is to quantify the mechanisms that control O₃ variability on monthly timescales and thus, that can modulate trends over limited time periods”*.

Also on **P3-4 L86-89** as follows: “It is important to highlight that to ensure statistical consistency, this study analyzes detrended anomalies rather than long-term trends. Therefore, the results describe variability-driven processes and should not be interpreted as a direct explanation of decadal O₃ trends caused by externally forced long-term changes in CO₂, N₂O emissions, or ODSs, but rather in the context of dynamical processes that are strongly influenced by the QBO”.

We also agree that some results were overinterpreted in the previous version of the manuscript, and we have corrected this. Since we no longer apply causal discovery to the subperiods, this issue does not arise in the revised manuscript.

1.4 Is it appropriate for this method to omit a temperature-mediated pathway between upwelling and ozone?

The paper explains some of the expert judgment required for formulating the DAG, including formulating a prior for the DAG and sometimes overruling linkages that are otherwise suggested by the algorithm. I am concerned about one such overruling, which is the direct connection between upwelling and ozone. At lines 178-9, it is stated; “Note, however, that we included one link assumption in this study, namely that the residual vertical velocity w^* does not directly influence O₃.” I assume that this overruling assumption is motivated by the fact that ozone around 10 hPa is in the photochemistry-dominated regime, and not the transport-dominated regime of the lower stratosphere. That is a good reason to block the advective connection between upwelling and ozone. However, I am concerned that this also blocks other connections between upwelling and ozone that are mediated by intermediate variables, in particular temperature. It seems plausible that an increase in upwelling can cool the stratosphere around 10 hPa, leading to an increase in ozone by changing the photochemical reaction rates. I am thinking both of the ozone source ($O + O_2 + M \rightarrow O_3 + M$, which speeds up at lower temperatures) and the ozone sinks (e.g., the Chapman sink, which is not the dominant sink but is very temperature sensitive, $O + O_3 \rightarrow 2O_2$).

As we now show in Fig. 5, for the 2004-2021 period, causal discovery robustly detects a bidirectional link from w^* to O₃ in the observation. This is now addressed in the manuscript on **P13 L283-287** as follows: “Causal discovery further detects a bidirected connection between w^* and O₃ in the observations, indicating the presence of a latent common driver of w^* and O₃ and that neither variable is an ancestor of the other (see Appendix B, Fig. B1a). As causal effect estimation requires a DAG, we define the direction from w^* to O₃ to quantify the strength of this link. This choice allows us to estimate the direct dynamical influence of w^* on O₃, which is sometimes identified by a causal discovery algorithm in sensitivity tests”.

In this study, the causal discovery algorithm does not use any link assumptions anymore.

In the QBO context, the omitted effects of temperature are not dominant compared to NO_x, but nor are they trivial. For example, Ming et al. (2025) found that NO_x explained about 65% of the ozone QBO at 10 hPa, with the remaining 35% due to other processes, which likely included temperature-mediated effects (and perhaps direct advection).

It would be helpful to see a discussion of this temperature-mediated connection. On the one hand, it could be correctly argued that it is not dominant. But nor is it necessarily negligible, so it would be helpful to know whether it contributes to noise in the fit, or, less ideally, whether it would modify the coefficients of the NO_x-mediated pathway. To be specific:

- In the NO_x-mediated pathway, strong upwelling leads to low NO_x and high O₃. Thus, this linkage in the DAG is negative.
- In the temperature-mediated pathway, strong upwelling leads to cooling that leads to higher O₃. This strong upwelling also separately leads to low NO_x. This would also lead to a linkage in the DAG from NO_x to O₃ that is negative.

Thus, the temperature-mediated pathway could be potentially amplifying the apparent strength of the NO_x-mediated pathway. Is this the case? If so, this would not necessarily be a major problem, depending on the intended use of the DAG, but it would be helpful if the consequences could be discussed, or, better yet, constrained. It would be a problem to the extent that this method was claimed to have produced a generic characterisation of the ozone sensitivity to upwelling perturbations, rather than a QBO-specific characterization.

We thank the Reviewer for this important point. To assess a temperature-mediated pathway in our framework, we performed additional causal discovery sensitivity tests in both the observations and the TOMCAT CTM simulation that included temperature (T) as an additional node in the causal graph alongside w^* , N₂O, NO₂, and O₃ in the tropical middle stratosphere. Unfortunately, as shown in Fig. 1 below, these tests did not yield a robust and physically interpretable structure: the inferred graphs frequently suggested a negative (contemporaneous and lagged) link from T to w^* and did not detect any influence on O₃. One of the reasons could be the diagnostic artifact, because in our analysis, w^* is not an independently observed quantity but a derived TEM variable that depends on thermodynamic fields (including T). Therefore, including both T and a temperature-dependent variable that represents transport in the same causal system violates the assumption that variables are measured independently, and it induces spurious directions and hides the interpretation of a temperature-mediated pathway.




Figure 1. Causal graphs detected by the causal discovery algorithm from the observations for the period 2004-2021 with different significance level α_{pc} and $\tau_{\max}=2$, when variable T is included.

An estimation of the temperature-mediated pathway would require reformulating the analyzed system so that the driver that represents stratospheric transport is independent of the temperature. This could be achieved by, for example, removing w^* and using N_2O as a proxy for the transport. But this would lead to a different graph structure, and it would no longer be directly comparable to the main analysis.

We now address this point on **P13 L287-294** as follows: “Based on Fig. 4, the direct influence of w^* on O_3 is much weaker, and a mediated pathway via N_2O and NO_2 dominates. However, temperature-mediated effects could amplify the apparent strength of the connection from NO_2 to O_3 , since enhanced upwelling induces both cooling (increasing O_3) and reduced NO_x species. To assess this, we performed additional analysis, including temperature anomalies in the tropical middle stratosphere. The inferred graph structure was not robust, likely because w^* is a derived diagnostic that depends on thermodynamic fields. Removing w^* altered the parent structure and prevented a direct comparison of direct causal effects shown in Fig. 4a. We therefore interpret the identified NO_2 -mediated pathway as the dominant mechanism, while acknowledging that temperature-related effects may potentially project onto this link”.

2 Specific comments

- Note that it is a bit confusing that Figure 1a shows data from 2002-2011 and then 2012-2018, but these are not actually the subperiods analyzed, because the first subperiod is 2004-2011 (due to the availability of N_2O data, Line 99).

Figure 1a is now removed.

- Is the end year of the subperiod inclusive or exclusive? If the end years are

exclusive, then is all the data from 2011 ignored, as visually suggested by Figure 1a?

Figure 1a is removed. It showed yearly anomalies (and probably confused even more), and therefore gave the impression that data from 2011 was ignored.

- Do the subperiods include the same number of months?

No, they did not. The refrain from using the subperiods in the revised manuscript.

- Line 33-34: A small point, but it is stated that trends in stratospheric ozone from 2000 to 2020 are “*unrelated to changes in ODS abundances*”. I think you might mean to say that ozone has declined despite ODS recovering trends, but these ODS trends could have still affected ozone (i.e., by reducing the magnitude of the decline).

As we have revised the Introduction, this discussion has been removed from the revised manuscript

- Line 67-68: “*Nonetheless, understanding the mechanisms that influence O3 behavior remains essential, as past trends do not preclude the possibility of their recurrence.*” This sentence is confusing. Do you mean that “trends that have occurred before can occur again”?

Now replaced with the following formulation to better align with the main scope of the paper on **P2 L44-46** as follows: “*However, the absence of a persistent long-term trend in recent analyses does not rule out future changes of O₃ in the tropical middle stratosphere. Similarly, negative trends observed in the early 2000s could recur if comparable dynamical–chemical conditions arise*”.

- Lines 228-229: The vertical shear of the zonal mean zonal wind is calculated as a proxy for the secondary upwelling circulation of the QBO, which is physically justified, but the shear is calculated over too narrow a vertical depth, only between 10 hPa and 12 hPa. Over this narrow a region, the shear is likely too noisy and would contain little information. I would recommend calculating the shear over a deeper layer, like 30 hPa to 10 hPa, to approximate the QBO vertical wind shear at 10 hPa. (Normally, I would recommend going above 10 hPa, but I see that 10 hPa is the top of the dataset, which is fine.)

We thank the Reviewer for this helpful suggestion, we revise Sect.3.4 and adopted our analysis accordingly.

- Line 245-246: “*It is important to note that w^{\wedge} does not exert a direct causal influence on NO_x, instead, its effect is mediated through N₂O.*” This implies that NO_x is not advected by upwelling. This would make sense if NO_x equilibrates photochemically much faster than transport can generate anomalies, but this

depends on whether one is thinking about NO_x or NO_y. Considering the long-lived chemical family of NO_y (see for example the discussion in Brasseur and Solomon (2005)), of which NO_x is a key constituent, transport can generate substantial anomalies. If NO_y were increased due to some external perturbation, then this would increase the amount of NO_y and this anomaly would be transported upwards with the flow, and it would bring along with it anomalies in NO_x as well. This allows for a direct pathway for upwelling to affect NO_x. This comment relates narrowly to the statement about the whether upwelling can affect NO_x and broadly to the formulation of the DAG, which does not include a direct arrow from upwelling to NO_x.

We now removed this statement from the revised manuscript. Also, we refrain from using link assumptions in the revised manuscript.

References

- Brasseur, G. P., and S. Solomon, 2005: *Aeronomy of the Middle Atmosphere: Chemistry and Physics of the Stratosphere and Mesosphere*. Springer, Dordrecht, Netherlands.
- Galytska, E., A. Rozanov, M. P. Chipperfield, S. S. Dhomse, M. Weber, C. Arosio, W. Feng, and J. P. Burrows, 2019: Dynamically controlled ozone decline in the tropical mid- stratosphere observed by SCIAMACHY. *Atmospheric Chemistry and Physics*, **19** (2), 767–783, doi:10.5194/ACP-19-767-2019.
- Ming, A., P. Hitchcock, C. Orbe, and K. Dubé, 2025: Phase and Amplitude Relationships Between Ozone, Temperature, and Circulation in the Quasi-Biennial Oscillation. *Journal of Geophysical Research: Atmospheres*, **130** (4), e2024JD042 469, doi: 10.1029/2024JD042469, e2024JD042469 2024JD042469.

Review comments on egusphere-2025-4745

We thank the Reviewer for providing feedback on this manuscript. We have now fully revised the manuscript. While the methodology remains the same, the manuscript now focuses on main chemical–dynamical pathway affecting tropical middle stratospheric ozone (O_3) over the single long period 2004–2021, instead of examining shorter time periods. The revised manuscript quantifies the processes that control O_3 variability on monthly timescales, which can also influence trends over shorter intervals. To better reflect the focus of the revised manuscript, we now propose an improved title: “*Causal inference for quantifying chemical–dynamical pathways controlling tropical middle stratospheric ozone variability*”. Below, we address the Reviewer’s comments one by one in blue. We use the notation **PL** to refer to the changes in the revised manuscript on a specific page and line.

This study designs a framework of causal inference to diagnose combined chemical and dynamical drivers of tropical middle-stratospheric ozone variability since 2004, highlighting differences between two subperiods of 2004-2011 and 2012-2018. Satellite observations and chemistry-transport simulations, as the input of the LPCMCP model, are used to explore a physical-chemical mechanism for ozone variability in this region. This work uses a broad tools methodologically and explains the causal framework technically, and the topic and framework are promising especially in the data-driven science. But I have several concerns about how the scientific motivation aligns with the dataset preprocessing, the interpretation of key links from the causal inference results. I would recommend a major revision because of the following concerns.

General comments

- The major motivation is contrasting “decline ozone (2004-2011)” and “increasing ozone (2012-2018)”. However, the analysis of causal inference process with two detrended subperiod, which removes low-frequency variability that may be part of the causal story. This leads to a mismatch between the aim of the study (interpretation of subperiod ozone changes) and what the method is actually applied to (detrended residual variation).

This has been fully revised. The analysis of two subperiods has been completely removed, and the focus is now on the full 2004–2021 period. The Introduction has been shortened and adapted to highlight the role of mechanisms controlling interannual ozone variability in the tropical middle stratosphere. We therefore removed from the manuscript Fig. 1a, which previously showed yearly ozone anomalies for the two subperiods. In the revised Introduction, we reduced the discussion of tropical ozone trends. However, the remaining discussion on tropical *middle* stratospheric ozone trends serves as the motivation to quantify the mechanism that governs interannual ozone variability.

On **P2 L49-50** we state: “While previously discussed O_3 trends motivate this study, our objective is to quantify the mechanisms that control O_3 variability on monthly timescales and thus, that can modulate trends over limited time periods”.

- **Line 283-285:** The paper’s one of the key findings is the 1 month lag for the path of $N_2O \rightarrow NO_2$ in 2012-2018, resulting from delayed NO_2 production due to the shorter N_2O residence time. One of my concern is that one month lag derives from monthly means dataset, which might be not reasonable. Because this lag could be caused by data preprocessing, like sampling, aggregation, or autocorrelation structure. Please add related arguments to clarify this point.

Since we do not analyze separate subperiods anymore, this small lag difference is no longer relevant and has been removed from the revised manuscript.

- **Line 171-172:** For the τ_{max} setting, you state that you ran sensitivity tests, which make the results more robust. But it ultimately use $\tau_{max} = 1$ in the discovery graphs, which limits the unfold of longer lag effect. Because the main narrative rely on lag interpretation, the main text should give more arguments on this point.

We thank the Reviewer for raising this point. Our revised manuscript now includes the sensitivity tests of the observed links to the choice of α_{pc} and τ_{max} for the analyzed period 2004-2021 in Fig. 5.

The graphs for the direct causal effect estimates for the observations and the TOMCAT CTM simulation shown in Fig. 4 are now based on $\alpha_{pc} = 0.05$ and $\tau_{max} = 2$. We would also like to highlight here that the detected observational graph is robust across tested α_{pc} and τ_{max} as shown in Fig. 5.

In the revised manuscript, we state the following on **P14 L308-318**: “To ensure the robustness of the detected in Fig.4a connections in observations in the tropical middle stratosphere during the period 2004-2021, Fig. 5 demonstrates the results of the application of the causal discovery algorithm with different setups of τ_{max} (depicted on the x-axis) and α_{pc} (depicted with markers). τ_{min} set to zero to account for the contemporaneous connections. It should be noted that choosing a τ_{max} that is too low risks missing causal links with longer delays, which violates the assumption of causal sufficiency. However, choosing τ_{max} too high without mitigation can dilute the detection power of the causal algorithm. Specifically, a larger τ_{max} expands the search, as the algorithm tests more possible lagged pairs. This leads to larger conditioning sets, which can further reduce the effect size and detection power. Therefore, it is important to condition only on a few relevant variables that actually explain the

relationship (Runge et al., 2019). Sensitivity tests from Fig. 5 show very similar results for different configurations of τ_{max} and α_{pc} . The chemical connections related to N_2O - NO_2 and NO_2 - O_3 pairs are robustly detected as contemporaneous across all experiments. The dynamical connections w^ - N_2O and w^* - O_3 are robustly detected with a one-month lag”.*

- **Line 294-298:** The paper states that LPCMCI does not robustly detect the expected chain in TOMCAT, and then uses the observation-derived graph as a fixed structure to estimate TOMCAT causal effects. Appendix C provides a useful diagnosis. This is a reasonable workflow. I confused by “given the observational causal template, does TOMCAT produce similar effect magnitudes?” Please highlight this point in the Result or Discussion.

We now add the discussion of the reasons why the causal discovery algorithm did not robustly detect the anticipated graph, in particular, the dynamical coupling of w^* with the N_2O . This is now addressed on **P13 L295-304** as follows: *“Unlike the observations, in the TOMCAT CTM simulation, the causal discovery algorithm does not fully reproduce the expected chemical-dynamical coupling as outlined in the Introduction, Sect. 4.1 and as shown in Fig. 4a. In particular, the anticipated one-month lagged link from w^* to N_2O is not robustly detected. We found that this occurs because O_3 exhibits very strong contemporaneous coupling with both N_2O and NO_2 , such that conditioning on O_3 makes the one-month lagged w^* to N_2O link statistically insignificant. As further demonstrated and discussed in Appendix B, this does not indicate a lack of dynamical coupling in the TOMCAT CTM simulation. Instead, it reflects the strong shared variability among the chemical tracers as the model uses a chemically consistent scheme for all the variables. To still assess the strength of the processes represented in TOMCAT, we did not rely on the TOMCAT-derived graph. Instead, by adopting a fixed graph structure derived from observational ground truth and expert knowledge, we can accurately estimate direct causal effects with the TOMCAT data (Fig.4b), providing a valid and pragmatic solution for quantifying model sensitivities”.*

Minor Comments

- 1) When you say sensitivity tests were done for α_{pc} and τ_{max} consider briefly summarizing the outcome in the main text.

Given the updated flow of the revised manuscript, we have now included the sensitivity test in the Sect. 4.2 of the manuscript, see Fig. 5. We therefore include the discussion of the robustness of the detected connections on **P14 L308-318** (as addressed above).

We would also like to note that the sensitivity analysis pointed to an interesting feature related to an emerging positive link from N_2O to NO_2 . On **P14 L319-323**, we discuss this as follows: “A further analysis of the sensitivity experiments reveals an additional feature in the N_2O - NO_2 relationship. A positive one-month lagged link from N_2O to NO_2 is detected across all tested τ_{max} but primarily at relaxed significance thresholds ($\alpha_{pc}=0.2$). For $\tau_{max}=1$, the link is also identified at $\alpha_{pc}=0.1$ and 0.05. Such a positive lagged connection is physically plausible, as NO_2 is produced from N_2O , as discussed in Table 1, and a delayed response may emerge at the monthly timeseries. However, given its sensitivity to the choice of α_{pc} this link cannot be considered a robust pathway”.