Second review of Galytska et al.: "Causal inference for quantifying chemical-dynamical pathways controlling tropical middle stratospheric ozone variability" submitted to ACP

April 8, 2026

The submitted paper has been overhauled to address the comments from my previous review, and it is now stronger, clearer, and better justified. The new Figure 1 is an excellent on-ramp to the content of the paper. Although the new material from the overhaul of the paper has raised some additional concerns, these concerns are not as fundamental, and I can envision a clear path to publication subject to Major Revisions.

I will begin by noting that the four major concerns from my previous review have all been satisfactorily addressed. Listed in order from my previous review: (#1) the new focus of the paper on variability rather than subperiod trends aligns the motivation and methods, (#2) there is now a discussion of the Simpson's paradox for N_2O, (#3) the apparent significance of small differences in lags has been removed, and (#4) the temperature-mediated pathway is analyzed and its apparent absence is appropriately discussed.

I have two Major comments on the overhauled manuscript:

- Lags with peak total effects: The manuscript explores the lagged temporal relationships between different variables, and argues that a lagging peak in the relationship of w* on NO2 is due to "sequential mediation" through intermediate variables. I think this explanation is problematic given the photochemical/dynamical context, and I propose an alternative.
- More clarity on causal methods: As this paper is translating causal modeling frameworks across disciplines, some more clarity on terms and the role of expert judgment is needed.

Below, I will expand upon my two major comments, and then discuss some minor comments.

1 Major comments

1.1 Lags with peak total effects

Lines 9-11: "The total causal effect (direct and mediated) peaks at a lag of approximately two-three months, indicating that the response develops on this timescale as the influence propagates through N2O and NO2."

Lines 372-374: "This peak defines the characteristic adjustment timescale of dynamical transport. The total effects of w* on NO2 (Fig. 7b) and on O3 (Fig. 7c and Appendix C) also reach their maxima at approximately three months, consistent with a sequential mediation through N2O and subsequently NO2."

The manuscript offers questionable explanations for the 3-month lag in the peak relationships between N2O and NO2. The paper refers to a dynamical adjustment timescale, but the paper is analyzing upwelling anomalies at the same altitude as the composition anomalies, so the effects of a spatiotemporally localized upwelling perturbation will be instantaneous. The paper also refers to sequential mediation, which I interpret as suggesting a perspective of thinking about the level of interest as a reaction chamber within which there are lags as reactants turn into products, such that a perturbation in the product lags behind a perturbation in the reactant.

I think this gives a misleading impression of the photochemical/dynamical context. Intuitively, this can be understood because the air at 10 hPa is continually moving upwards, so an anomaly in composition at 10 hPa will be transported up and away, and the composition at 10 hPa will then depend on the anomalies that are being transported up from below. A hypothetical pulse in upwelling that is highly localized in space (at 10 hPa) and time (to much less than one month) will lead to an instantaneous response in composition in that month, but its effects will be largely transported up and away from 10 hPa by the next month. Lagged responses therefore result from the spatiotemporal structure of the pulse and covariations in upwelling at lower altitudes with what is happening at 10 hPa.

(There is also little evidence that the photochemical timescales would lead to "sequential mediation" after 3 months. At 10 hPa, N2O has a photochemical timescale of about 3 years (Fig. 5.32, Brasseur and Solomon, 2005), NOx has a photochemical timescale of about 10 days (Fig. 5.37, Brasseur and Solomon, 2005), and odd oxygen has a timescale of about 10 days. None of these suggest sequential mediation that leads to a peaked 3-month response.)

I suspect that the temporal lag reflects the spatiotemporal structure of the covariability between upwelling at 10 hPa and upwelling at lower altitudes. Such anomalies in upwelling are primarily dominated by the QBO, and these anomalies in w* descend along with the shear zones. This means that an upwelling anomaly first experienced at 10 hPa will have a modest instantaneous and local effect on composition, but after some months, the upwelling anomaly will still be present at 10 hPa but will also be affecting lower altitudes, leading to a cumulatively larger effect on N2O and NO2, and therefore O3. The lag at which total effect reaches a maximum would therefore depend on the depth of the upwelling anomaly,

and the rate at which QBO shear zones descend. If the QBO shear zone has a characteristic depth of 5 km and the QBO has a characteristic descent rate of 1 km/month, this would suggest a lag timescale of roughly 5 months. I suspect the local maximum in the lag in Fig 7 at 3 months is not significantly distinguished from the lag at 2 months or 4 months.

If the authors believe my suggested interpretation is misguided, then I would be consider a justification for their interpretation against this plausible alternative.

As a general comment that builds on the above considerations, please discuss the limitations of your framework as compared to a spatiotemporal model of the system. The framework here considers only the behavior at a single level, but this level is in the tropical pipe where air is continually ascending, an important piece of dynamical context enters subtly. It suggests that local reaction lags (in the spirit of sequentially mediated reactions in a chamber) might be conceptually hard to distinguish from the response to correlated changes in dynamical variables at lower altitudes.

1.2 More clarity on causal methods

This manuscript will be introducing many concepts from causal modeling that will be new to the majority subset of readers from atmospheric science. The paper is generally quite elegant at bridging these communities, but I think some more clarity could be achieved with consistency of language and identification of methods. I have a couple comments that fall broadly under this heading.

This is optional, but I think it would be very helpful towards the spirit of this comment: some of the causal language in the paper is not defined, which risks that the reader will fall back on their colloquial interpretation of these terms, e.g., mediator. It would be helpful to have something resembling a glossary for some of the causal terms used in this paper. It could be in the form a table, and it could connect colloquial understanding of causality in the atmospheric sciences to technical meanings in this paper. It could clarify which causal terms are being used in a technical sense versus in a colloquial sense. Some of the terms to include are:

- "causal effect"
- "total causal effect"
- "direct causal effect"
- "dependency"
- "connection"
- "mediator"
- "latent"

- "direct effect" (is this the same as direct causal effect? If so, probably best to use the full term each time or note the elision explicitly)
- "total effect" (is this the same as total causal effect? If so, probably best to use the full term each time or note the elision explicitly)

This is optional, because I don't want to set an unrealistic standard of generality, but I hope the authors will consider the pros and cons. I have regularly gotten lost in the paper reading about total vs. direct effects and not having a clear sense of the distinction.

Along these lines, you could comment on how "total causal effects" and "direct causal effects" relate specifically to linear regression. Is there a special case in which they would be equal to each other for a sufficiently simple causal network? How should readers interpret the magnitude of "causal effects", i.e., what does it mean to have a causal effect of -1 or 1? Can it exceed 1? Is it analogous to a regression slope, or r, or r^2, or something else entirely?

Please comment briefly in the paper on how Wright's method relates to more modern Bayesian structural causal modeling. My sense from reading Judah Pearl's *Book of Why* is that there has been a revolution in causal inference, and this postdates Wright's methods. Compared to Pearl's approaches, are Wright's methods more fundamental, analogous, unrelated, identical, a simple case of a more complex new methodological landscape ... ?

In the response to reviewers, the authors wrote: "In this study, the causal discovery algorithm does not use any link assumptions anymore." However, expert judgment is still discussed as a part of the method, and it is included in Fig 2. How does expert judgment enter at this point. Is it through "triangulation" (Line 265)? Please define triangulation in this context, and explain how you specifically used it.

Lines 182-185: What form of regression is used for the links in Wright's method? is it linear regression on each parent separately? is it multiple linear regression on all parents simultaneously?

Line 166: "α_{pc} and τ_{max}": I believe in the latest version of the manuscript these are not defined. Please define them when first used.

2 Minor comments about the Introduction

Lines 22-24: "A balance between photochemical production and loss mainly determines the overall abundance of stratospheric O3. Meanwhile, its global distribution and inter-annual changes are mainly determined by dynamical and chemical processes..."

This introductory comment suggests a dichotomy where photochemistry determines the total amount of ozone and transport determines the distribution. I know this dichotomy is meant to be notional, but it is inaccurate. Ozone in the tropical lower stratosphere is damped by transport (e.g., Brasseur and Solomon 2005, section 3.5.2), and this affects its overall abundance as well as the fact that ozone has an interior maximum. Match

et al. (2025) argued that the shape of the tropical ozone layer can only be understood as a photochemical-transport equilibrium.

Lines 76-77: "The positive relationship between N2O and O3 (panel d) is consistent with both tracers exhibiting lifetimes that exceed their vertical transport timescales in this region (Bnisch et al., 2011)."

This is not correct for ozone. The photochemical lifetime of odd oxygen at 30 km is shorter than the transport timescale. Ozone at this level is understood to be under photochemical control and not dynamical control, e.g., Brasseur and Solomon, 2005, Fig 5.11

Line 78: "While the chemical-dynamical coupling governing tropical middle-stratospheric O3 is understood"

In my view, the coupling problem remains incompletely understood. For example, it was recently shown in Hitchcock and Ming (2025) that ozone can play a leading order in modulating the QBO secondary circulation, which had not been fully appreciated, and the APARC QUOCA activity is seeking to better understand this chemical-dynamical coupling of the QBO (Orbe et al., 2025). A better qualified statement would be that the key chemical reactions and dynamical variables have been identified and can be simulated in chemistry-climate models.

Line 93: "thermodynamical regimes" I am not sure if you mean to emphasize the temperature QBO here by saying thermodynamical, but the dynamical regime is likely more relevant for the transport processes considered in this paper.

References

Hitchcock, P., and A. Ming, 2025: The Role of Ozone in the Secondary Circulation of the QBO: Linear Theory. *Journal of Geophysical Research: Atmospheres*, **130 (20)**, e2025JD044 766, doi:10.1029/2025JD044766, e2025JD044766 2025JD044766.

Match, A., E. P. Gerber, and S. Fueglistaler, 2025: Protection without poison: Why tropical ozone maximizes in the interior of the atmosphere. *Atmospheric Chemistry and Physics*, **25 (8)**, 4349–4366, doi:10.5194/acp-25-4349-2025.

Orbe, C., and Coauthors, 2025: Experimental Protocol for Phase 1 of the APARC QUOCA (QUasibiennial oscillation and Ozone Chemistry interactions in the Atmosphere) Working Group. *EGUsphere*, 1–36, doi:10.5194/egusphere-2025-2761.