Title: "Regressing Away Common Neural Choice Signals does not make them Artifacts. Comment on Frömer et al (2024, Nature Human Behaviour)"

Authors: Redmond G. O'Connell[1], Elaine A. Corbett[2], Elisabeth Parés-Pujolràs[2], Daniel Feuerriegel[3] & Simon P. Kelly[2]

1. Trinity College Institute of Neuroscience and School of Psychology, Trinity College Dublin, Dublin, Ireland 2. School of Electrical and Electronic Engineering and UCD Centre for Biomedical Engineering, University College Dublin, Dublin, Ireland. 3. Melbourne School of Psychological Sciences, University of Melbourne, Melbourne, Australia
E-mail: reoconne@tcd.ie, simon.kelly@ucd.ie

Acknowledgments: This work was funded by an SFI research grant to R.G.O. and S.P.K. (19/US/3599). R.G.O. was funded also by the European Research Council Consolidator Grant IndDecision – 865474, and S.P.K. also by a Wellcome Trust Investigator Award (219572/Z/19/Z). D.F. was supported by an Australian Research Council Discovery Early Career Researcher Award (ARC DE220101508). E.A.C. was funded by the European Research Council Starting Grant Myodecision (101077772).

Author Contributions: R.G.O. & S.P.K. drafted the reply. S.P.K. and E.P. performed simulations. All authors contributed text and approved the final version.

Data Availability: Simulation code can be found here:
https://github.com/Elisabeth-Pares/NHB_MA_OConnell_etal_2024

Abstract:

Frömer et al (2024, Nature Human Behaviour) report that application of a deconvolution method to correct for component overlap in the event-related potential eliminates signatures of sensory evidence accumulation from response-aligned grand average measurements of the centro-parietal positivity (CPP), suggesting that these signatures arise artifactually. Here, we argue that the analysis and interpretation of their perceptual choice data are critically flawed. We demonstrate with simulations that the deconvolution analyses used by the authors are not designed to correctly capture bounded accumulation signals. Applying these same analyses to true, ramping, accumulation-to-bound signals more closely mimics Frömer et al's empirical results than the alternative scenario in which such signals are absent.

Main Text:

Frömer et al (2024) describes an experiment devised to identify neural correlates of value-based choice in the human event-related potential (ERP). Their analyses focus on a component known as the centro-parietal positivity (CPP) that had been widely implicated in tracing the sensory evidence accumulation (EA) process underpinning perceptual decisions (O'Connell & Kelly, 2021) but rarely examined in the context of value-based decisions. Frömer et al's paper makes a valuable contribution in providing compelling evidence, from response-locked waveforms divided by response time, that the CPP does not trace the value-based choice process that culminates in action selection.

In addition to their analysis of the value-based choice data, the authors analyzed ERPs from three pre-existing perceptual choice datasets. In these datasets, the response-aligned, trial-averaged CPP waveforms exhibit amplitude modulations as a function of evidence strength or reaction time (RT) that have been widely reported in previous research and interpreted as empirical support for the CPP's role in sensory EA. However, when a deconvolution method was applied, which models the ERP as the sum of stimulus-locked (S) and response-locked (R) components that overlap according to RT, these pre-response effects no longer reach statistical significance. From this the authors conclude that the CPP is an entirely stimulus-locked component, and that its pre-response amplitude modulations are an artifact of averaging the overlapping S and R components across trials. They go on to suggest that, rather than tracing sensory EA, the CPP may be linked to various alternative processes such as appraisal or attention enhancement in value-based contexts, and stimulus onset detection or performance monitoring in perceptual contexts.

Here, we argue that the analysis and interpretation of these perceptual choice data are critically flawed. We first demonstrate with simulations that the deconvolution analyses used by the authors are not designed to correctly capture bounded accumulation signals. Applying these analyses to true, ramping, accumulation-to-bound signals more closely mimics Frömer et al's empirical results than the alternative scenario in which such signals are absent. Second, we highlight key empirical results that have been misinterpreted by the authors. Finally, since the authors offer only a partial picture of the previous literature identifying the CPP with sensory EA, we list numerous signatures of EA that have been identified in the CPP, other than the response-aligned average waveform criteria on which the present work narrowly focused.

<u>Unfold misinterprets evidence accumulation signals</u>

The Unfold toolbox (Ehinger & Dimigen, 2019) implements deconvolution (i.e. overlap-correction) under the assumption that the only signals present in the EEG are neural responses precisely time-locked to either the stimulus or response (Fig 1a). Indeed, in this situation the RT-dependent summation of the S and R components can cause false RT effects in the average response-locked data even when there is no RT effect in the underlying ground-truth components, and this ground-truth is recovered after applying Unfold (Fig 1b&c; recapitulating Frömer et al Figure 5). However, the accumulation-to-bound processes invoked in sequential sampling models are fundamentally different. In standard diffusion models featuring accumulation of stationary evidence, the accumulation process begins at (or shortly after) the stimulus and ends at response. In this way it is neither stimulus-locked nor response-locked, but rather interposed between the two, temporally stretching/contracting in slow-/fast-RT trials (Fig 1d&e). There is currently no means to capture such timescale-variation directly in the Unfold toolbox. To demonstrate the impact of this, we applied the same Unfold analyses as were applied to the contrast discrimination dataset (Steinemann et al., 2018), to a simulated S-R-interposed ramp signal. This task requires an immediately-reported decision about a single, stationary sensory feature (contrast difference), whose accumulation would approximate a straightforward, S-R-interposed ramp.

As expected, the Unfold analysis spuriously ascribes an initial portion of the S-R-interposed ramp to the S component and the end portion to the R component (Fig 1f). In so doing, Unfold systematically blunts RT effects that are truly present in the data, because the inappropriate subtraction of a spurious S component reduces the fast-RT waveform close to response while reducing the slow-RT waveform further back in time. Thus, while Unfold does correctly eliminate RT effects where they falsely arise from S- and R-component overlap, it also erroneously removes RT effects when they are truly present in a ramping accumulator signal. Frömer et al performed two additional Unfold analyses on the contrast discrimination dataset: one in which both the S and R components were allowed to vary as a function of RT (their Supplementary Figure 6C) and another in which neither component varied with RT (Supplementary Figure 7). Running these analyses on the S-R interposed ramp in our simulations produced the same qualitative patterns observed in the empirical data, key elements of which were not observed when applying them to RT-invariant S and R components (Figure 1g-h).

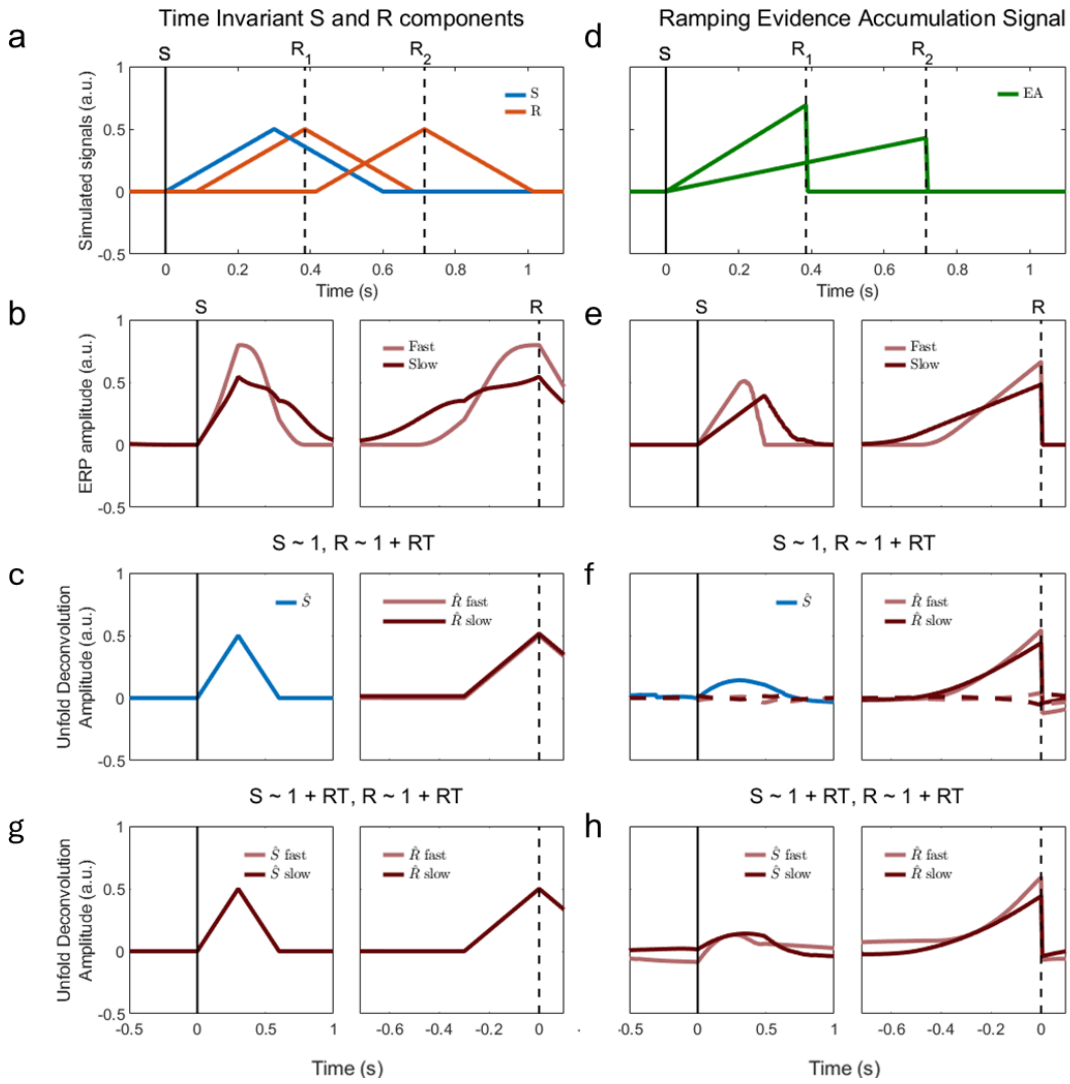

Figure 1: *Outcomes of using Unfold on simulated data that do (timescale-invariant S and R components), versus do not (S-R interposed ramp), adhere to the underlying assumptions of Unfold. We first generated ground-truth simulated data with only timescale-invariant S and R components, and no effect of RT on the amplitudes of either the S or R component.* ***A)*** *Two example trials from this scenario are shown, one around*

the 10th (R_1, faster) and one around the 90th (R_2, slower) RT percentile. ***B)*** *Trial-averaged timecourses aligned to stimulus and response, showing that RT-dependent overlap causes RT effects in the response-locked traces.* ***C)*** *"Correcting" the signal by removing the estimated S-component (blue) from each trial leaves a response-aligned waveform that precisely matches the R-component created in the ground truth, with no RT effects.* ***D)*** *In an alternative simulated scenario, there are no S or R components but rather a single accumulation-to-bound process ramping from stimulus to response, with a collapsing bound.* ***E)*** *In this ground truth, there are effects of RT on both the buildup rate and amplitude reached at response when average traces are split by median RT (e.g. as observed in (Steinemann et al., 2018 & Kelly et al., 2021).* ***F)*** *Running the signals through Unfold in the same way as for panel C, however, spuriously allocates some of the initial part of the ramp to the S component (blue) whose subtraction dramatically reduces the real RT effects on both slope and amplitude in the response-aligned trace, as observed in Frömer et al (see their Figure 6). This spurious reduction in RT-dependent buildup occurs for constant bounds also; we simulated a collapsing bound here to produce qualitatively similar patterns as in the 'uncorrected' contrast discrimination data of Steinemann et al. Frömer et al went on to conduct two more versions of the Unfold analysis to elaborate their claims. First, they sought to test whether, even if reduced from the response-locked waveforms, the RT effects may appear instead in the stimulus-locked waveforms after removing the R components. Whereas in their main analysis only the R component was allowed to vary between fast and slow RT categories, here they allowed RT effects in both the S and R components. They found that stimulus-locked waveforms for fast- and slow-RT trials had similar peak amplitudes, but with a large RT effect (fast < slow) now appearing in the baseline period (their Supplementary Figure 6C).* ***H)*** *This is in fact exactly what occurs when the same analysis is applied to the ramp signal ground truth. Note that the degree and morphology of such artifacts arising from Unfold depend strongly on analysis settings (e.g. time-expansion window, regression structure, etc), and while we have attempted to match those applied to the Steinemann et al dataset here, the code that was shared with the paper indicates that other perceptual decision datasets were analysed with different settings.* ***G)*** *In the RT-invariant S/R component ground truth, these baseline shifts do not occur indicating that they arise due to the Unfold model not having any regressor included which can capture the essential characteristics of an S-R-interposed ramp. Second, Frömer et al reported that removing both the R and S components of an RT-agnostic regression leaves very little residual activity in the empirical ERP data (see their Supplementary Figure 7), which they interpret as evidence that Unfold has not removed real EA activity. Again however, we observe qualitatively the same outcome when this subtraction is performed on the simulated ramp data (dashed traces in panel F). Thus, systematic misinterpretations arise from the fact that these Unfold analyses do not directly test for the presence or absence of an accumulation signal. In order to truly do so, the deconvolution model would need to have predictors that are able to capture a ramp whose timescale varies with RT through including a ramp regressor in the design matrix; we will examine this in detail on real and simulated data in an article in preparation.*

<u>Signatures of evidence accumulation remain evident after deconvolution with Unfold</u>

For ease of illustration, the above simulation involved a contrived, minimal ramp shape. In more fully-elaborated EA models commonly fit to behavioral data, additional 'non-decision' delays are assumed to occur before and after the noisy accumulation process. Simulating this using code that we provided, Frömer et al show that the more latency jitter in the post-accumulation (motor) delay relative to pre-accumulation (stimulus encoding/transmission) delay, the more a true EA signal is attributed to the S component, and even in this case, its amplitude is larger for fast compared to slow RTs (right panel Supplementary Figure 5).

From this, the authors conclude that, if Unfold has in fact attributed EA activity to the S component of the real EEG data, then it should exhibit a similar amplitude modulation.

When the authors tested this, of the two datasets that exhibited the relevant amplitude modulation in their uncorrected stimulus-locked waveforms (visual search and contrast discrimination), both also showed those modulations in their S components after deconvolution (Supplementary Figure 6b&c). In the case of the contrast discrimination data, deconvolution causes the RT effect on amplitude to transfer to the pre-stimulus baseline period, so that the peak amplitude effect would emerge with baseline-correction, as noted by Frömer et al. As our Figure 1h shows, the same spurious pre-stimulus amplitude difference arises as an expected artifact of applying Unfold to a ramping signal, but not for the overlapping stimulus- and response-locked components assumed by Frömer et al (Figure 1g). In this way, the Unfold results of Frömer et al actually constitute positive evidence that the observed signals arise from a ramping EA signal.

We note however that the simulations in their Supplementary Figures 5 shows that it is not possible to make clear predictions regarding where EA signatures should manifest after applying Unfold without first knowing the relative amount of pre- versus post-EA latency jitter. For example, in the case where pre-EA latency variability was set to double that of post-EA variability (left panel), the amplitude modulations by RT are more evenly split between the S and R components, rendering them more difficult to detect in statistical tests applied to one component at a time. This reflects the broader fact that the degree to which any EA signal will be aligned to the final response will be highly task-dependent. For example, pre-EA variability can be expected to increase in tasks in which the location of the physical evidence is not known in advance while post-EA variability will increase with the complexity and speed of the decision-reporting actions.

<u>Various signatures of EA have been identified in the CPP without recourse to analyses of response-aligned average waveforms</u>

Whereas Frömer et al focus their analyses and discussion of the previous literature on amplitude modulations in trial-averaged response-locked waveforms, this overlooks the fact that many previous studies have employed hypothesis-driven analyses and experimental manipulations to establish a range of other signatures of sensory EA in the CPP. These include, but are not limited to: single-trial surface plots that exhibit a progressive build up extending from shortly after evidence onset until the perceptual choice report (Simon et al., 2018; Loughnane et al., 2016; O'Connell et al., 2012); after statistically controlling for RT, single-trial measurements of the CPP's pre-response amplitude are modulated by experimental manipulations of time pressure in a manner that accords with the decision bound adjustments identified by sequential sampling models (Kelly et al., 2021; Steinemann et al., 2018); the pre-response amplitude of the CPP is larger on trials in which the correct alternative was invalidly cued despite slower RTs also in accordance with model-identified boundary adjustments (Kelly et al., 2021; Kohl et al., 2020); subtle perturbations (O'Connell et al., 2012) or random fluctuations (Devine et al., 2019) of the physical evidence while decision formation is ongoing leads to corresponding variations in the CPP's build-up in stimulus-aligned measurements that predict subsequent behavioral effects; both stimulus- and response-aligned single-trial measurements of the CPP's build-up rate are correlated with choice confidence after controlling for RT, accuracy and difficulty (Dou et al., 2024); the CPP's stimulus and response-aligned average timecourse as well as modulations of its pre-response amplitude as a function of RT, evidence strength, accuracy, time pressure and prior knowledge can all be recapitulated by simulating cumulative evidence signals from a sequential sampling model that provides excellent behavioral fits (Kelly et al., 2021).

If considered only one at a time, some of the above results could potentially be explained by invoking mechanisms other than sensory EA. However, sensory EA is the only account that we are currently aware of

that can accommodate all of the above observations at once and in detail. Finally we note that the CPP features observed even in the value-based choice data of Frömer et al are also consistent with it tracing the perceptual decisions (i.e. stimulus identification) that are a prerequisite for initiating value comparisons on this task. These features include the fact that its peak falls well before the final response (the value choice process can only start once the stimuli have been identified), that its peak falls closer to the final response on trials with faster RTs (faster value choices will cause the response-locked ERP to overlap more with the perceptual choice process), that Unfold assigns it to the S component and that, once isolated in the S component, it no longer exhibits any value choice-related modulations. While the above observations accord with the CPP tracing perceptual decisions for the value task, an interesting challenge for future research will be to determine whether those decisions take the form of an EA process.

In summary, we show that the regression-based framework used in Frömer et al. (2024) is unsuitable for correctly identifying and characterizing EA signals. There is also extensive prior evidence that the CPP traces sensory EA, and this evidence is not dependent on analyses of response-locked ERPs. Thus, while the authors make a valuable contribution in showing CPP traces during value comparisons that suggest it does not participate in forging the ultimate value-based choice, their deconvolution-based claims relating to the CPP and sensory EA are not supported by their data. Component overlap is undoubtedly an important issue to account for in noninvasive neurophysiology, and to the extent that paradigm design cannot fully eradicate them (Kelly & O'Connell, 2015), analysis tools such as deconvolution can in principle be helpful; however, as the above discussions highlight, further work is needed to develop deconvolution methods that are equipped to detect the presence or absence of EA signals.

References:

Devine, C. A., Gaffney, C., Loughnane, G., Kelly, S. P., & O'Connell, R. G. (2019). The Role of Premature Evidence Accumulation in Making Difficult Perceptual Decisions under Temporal Uncertainty. *eLIfe* 27:8:e48526

Dou, W., Martinez Arango, L. J., Castaneda, O. G., Arellano, L., Mcintyre, E., Yballa, C., & Samaha, J. (2024). Neural Signatures of Evidence Accumulation Encode Subjective Perceptual Confidence Independent of Performance. *Psychological Science*, *35*(7), 760–779.

Ehinger, B. V., & Dimigen, O. (2019). Unfold: an integrated toolbox for overlap correction, non-linear modeling, and regression-based EEG analysis. *PeerJ*, *7*, e7838.

Frömer, R., Nassar, M. R., Ehinger, B. V., & Shenhav, A. (2024). Common neural choice signals can emerge artefactually amid multiple distinct value signals. *Nature Human Behaviour*. https://doi.org/10.1038/s41562-024-01971-z

Kelly, S. P., Corbett, E. A., & O'Connell, R. G. (2021). Neurocomputational mechanisms of prior-informed perceptual decision-making in humans. *Nature Human Behaviour*, *5*(4), 467–481.

Kelly, S. P., & O'Connell, R. G. (2015). The neural processes underlying perceptual decision making in humans: recent progress and future directions. *Journal of Physiology, Paris*, *109*(1-3), 27–37.

Kohl, C., Spieser, L., Forster, B., Bestmann, S., & Yarrow, K. (2020). Centroparietal activity mirrors the decision variable when tracking biased and time-varying sensory evidence. *Cognitive Psychology*, *122*, 101321.

Loughnane, G. M., Newman, D. P., Bellgrove, M. A., Lalor, E. C., Kelly, S. P., & O'Connell, R. G. (2016). Target Selection Signals Influence Perceptual Decisions by Modulating the Onset and Rate of Evidence Accumulation. *Current Biology: CB*, *26*(4), 496–502.

O'Connell, R. G., Dockree, P. M., & Kelly, S. P. (2012). A supramodal accumulation-to-bound signal that determines perceptual decisions in humans. *Nature Neuroscience*, *15*(12), 1729–1735.

O'Connell, R. G., & Kelly, S. P. (2021). Neurophysiology of Human Perceptual Decision-Making. *Annual Review of Neuroscience*, *44*, 495–516.

Simon, D. M., Nidiffer, A. R., & Wallace, M. T. (2018). Single Trial Plasticity in Evidence Accumulation Underlies Rapid Recalibration to Asynchronous Audiovisual Speech. *Scientific Reports*, *8*(1), 12499.

Steinemann, N. A., O'Connell, R. G., & Kelly, S. P. (2018). Decisions are expedited through multiple neural adjustments spanning the sensorimotor hierarchy. *Nature Communications*, *9*(1), 3627.