Accepted Manuscript

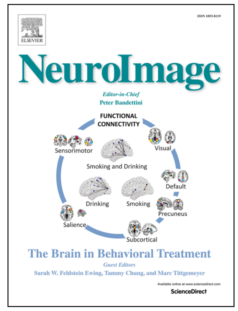

Solved problems for Granger causality in neuroscience: A response to Stokes and Purdon

Lionel Barnett, Adam B. Barrett, Anil K. Seth

PII: S1053-8119(18)30493-2

DOI: 10.1016/j.neuroimage.2018.05.067

Reference: YNIMG 14992

To appear in: *NeuroImage*

Received Date: 5 February 2018

Revised Date: 1 May 2018

Accepted Date: 27 May 2018

Please cite this article as: Barnett, L., Barrett, A.B., Seth, A.K., Solved problems for Granger causality in neuroscience: A response to Stokes and Purdon, *NeuroImage* (2018), doi: 10.1016/j.neuroimage.2018.05.067.

Solved problems for Granger causality in neuroscience: A response to Stokes and Purdon

Lionel Barnett, Adam B. Barrett, Anil K. Seth*

Sackler Centre for Consciousness Science and Department of Informatics,
University of Sussex, Brighton BN1 9QJ, UK
**a.k.seth@sussex.ac.uk (correspondence)*

Abstract

Granger-Geweke causality (GGC) is a powerful and popular method for identifying directed functional ('causal') connectivity in neuroscience. In a recent paper, Stokes and Purdon (2017b) raise several concerns about its use. They make two primary claims: (1) that GGC estimates may be severely biased or of high variance, and (2) that GGC fails to reveal the full structural/causal mechanisms of a system. However, these claims rest, respectively, on an incomplete evaluation of the literature, and a misconception about what GGC can be said to measure. Here we explain how existing approaches resolve the first issue, and discuss the frequently-misunderstood distinction between *functional* and *effective* neural connectivity which underlies Stokes and Purdon's second claim.

Keywords: Granger causality, functional connectivity, effective connectivity, statistical inference

Granger-Geweke causality (GGC) is a powerful analysis method for inferring directed functional ('causal') connectivity from time-series data, which has become increasingly popular in a variety of neuroimaging contexts (Hesse et al., 2003; Roebroeck et al., 2005; Ding et al., 2006; Dhamala et al., 2008a; Bressler and Seth, 2011; Valdes-Sosa et al., 2011; Barrett et al., 2012; Seth et al., 2015). GGC operationalises a statistical, predictive notion of causality in which causes precede, and help predict their effects. When implemented using autoregressive modelling, GGC can be computed in both time and frequency domains, in both bivariate and multivariate (conditional) formulations. Despite its popularity and power, the use of GGC in neuroscience and neuroimaging has remained controversial. In a recent paper, Stokes and Purdon (2017b) raise two primary concerns: (1) that GGC estimates may be severely biased or of high variance, and (2) that GGC fails to reveal the full structural/causal mechanisms of a system. We explain why these concerns are misplaced.

We note that Stokes and Purdon (2017a) have since responded to critiques of their claims by Barnett et al. (2017)[1] and Faes et al. (2017). Here, we expand on the points made in those articles [see also Dhamala et al. (2018)], and reply in detail to Stokes and Purdon (2017a).

Regarding the first claim, Stokes and Purdon (2017b) describe how bias and variance in GGC estimation arise from the use of separate, independent, full and reduced regressions. While true, this problem has long been recognised (Chen et al., 2006; Barnett and Seth, 2014), and has already been solved by methods which derive GGC from a single full regression[2]. These methods essentially extract reduced model parameters from the full model via factorisation of the spectral density matrix. Well-documented approaches include Wilson's frequency-domain algorithm (Wilson, 1972; Dhamala et al., 2008b, 2018), Whittle's time-domain algorithm (Whittle, 1963; Barnett and Seth, 2014), and a state-space approach which devolves to solution of a discrete-time algebraic Riccati equation (Lancaster and Rodman, 1995; Barnett and Seth, 2015; Solo, 2016). Thus, the source of bias and variance discussed in Stokes and Purdon (2017b) has already been addressed and resolved by previously published methods.

In their reply, Stokes and Purdon (2017a) acknowledge some of this work by saying: "We also described the state space solution to these problems in Dr. Stokes' Ph.D. thesis [Stokes (2015)] in January 2015, but felt it was important to first characterize and describe the problem, before laying out a solution to that problem." It is however worth noting that, at that time, the problem itself was already long acknowledged (Chen et al., 2006) and, even prior to publication of the state-space method, the distinct and equally effective methods of Dhamala et al. (2008b) and Barnett and Seth (2014) were already in the public domain.

To further illustrate the issue of bias and variance highlighted by Stokes and Purdon (2017a), and its resolution by single-regression methods, in Fig. 1 we plot estimated frequency-domain GGC for the 3-node vector-autoregressive (VAR) model in Stokes and Purdon (2017b), Example 1, using the single-regression state-space method (Barnett and Seth, 2015; Solo, 2016); see also Faes et al. (2017), Figure 1 and Dhamala et al. (2018), Fig. 1. We remark that identical results are obtained using the time-domain spectral factorisation method of Barnett and Seth (2014), as implemented in the current (v1.0, 2012) release of the associated MVGC Matlab© software package (Barnett and Seth, 2012). Our Fig. 1 may be directly compared with Fig. 2 in Stokes and Purdon (2017b);

[1] Barnett et al. (2017) is a preprint of an earlier version of the current article.
[2] But note that the "partition matrix" solution proposed by Chen et al. (2006) is incorrect; see, e.g., Solo (2016).