Journal of Econometrics 192 (2016) 207–230

Contents lists available at ScienceDirect

Journal of Econometrics

journal homepage: www.elsevier.com/locate/jeconom

Testing for Granger causality with mixed frequency data[☆]

Eric Ghysels[a,*], Jonathan B. Hill[b], Kaiji Motegi[c]

[a] Department of Economics and Department of Finance, Kenan-Flagler Business School, University of North Carolina at Chapel Hill, United States
[b] Department of Economics, UNC Chapel Hill, United States
[c] Faculty of Political Science and Economics, Waseda University, Japan

ARTICLE INFO

Article history:
Received 13 March 2014
Received in revised form
1 July 2015
Accepted 9 July 2015
Available online 29 December 2015

Keywords:
Granger causality test
Local asymptotic power
Mixed data sampling (MIDAS)
Temporal aggregation
Vector autoregression (VAR)

ABSTRACT

We develop Granger causality tests that apply directly to data sampled at different frequencies. We show that taking advantage of mixed frequency data allows us to better recover causal relationships when compared to the conventional common low frequency approach. We also show that the new causality tests have higher local asymptotic power as well as more power in finite samples compared to conventional tests. In an empirical application involving U.S. macroeconomic indicators, we show that the mixed frequency approach and the low frequency approach produce very different causal implications, with the former yielding more intuitively appealing result.

1. Introduction

It is well known that temporal aggregation may have spurious effects on testing for Granger causality, as noted by Clive Granger himself in a number of papers, see e.g. Granger (1980, 1988) and Granger and Lin (1995).[1] It is worth noting that whenever Granger causality and temporal aggregation are discussed, it is typically done in a setting where *all* series are subject to temporal aggregation. In such a setting it is well-known that even the simplest models, like a bivariate VAR(1) with stock (or skipped) sampling, may suffer from spuriously hidden or generated causality, and recovering the original causal pattern is very hard or even impossible in general.

In this paper we deal with what might be an obvious, yet largely overlooked remedy. Time series processes are often sampled at different frequencies and then typically aggregated to the common lowest frequency to test for Granger causality. The analysis of the present paper pertains to comparing testing for Granger causality with all series aggregated to a common low frequency, and testing for Granger causality taking advantage of all the series sampled at whatever frequency they are available. We rely on mixed frequency vector autoregressive, henceforth MF-VAR, models to implement a new class of Granger causality tests.

We show that mixed frequency Granger causality tests better recover causal patterns in an underlying high frequency process compared to the traditional low frequency, henceforth LF, approach. We also formally prove that mixed frequency, henceforth MF, causality tests have higher asymptotic power against local alternatives and show via simulation that this also holds in finite samples involving realistic data generating processes. The simulations indicate that the MF-VAR approach works well for small differences in sampling frequencies—like quarterly/monthly mixtures.

We apply the MF causality tests to monthly U.S. inflation, monthly crude oil price fluctuations and quarterly real GDP growth. We also apply the conventional causality test to the aggregated quarterly price series and real GDP for comparison. These two approaches yield very different conclusions regarding causal patterns. In particular, significant causality from oil prices

☆ The first author acknowledges support of a Marie Curie FP7-PEOPLE-2010-IIF grant and benefited from funding by the Federal Reserve Bank of New York through the Resident Scholar Program as well as the European Central Bank through the Duisenberg Scholar Program. We thank participants at 2014 Japanese Joint Statistical Meeting (University of Tokyo, September) for helpful comments. We also thank two referees, an associate editor, and co-editor Oliver Linton for helpful comments.

* Corresponding author.
E-mail addresses: eghysels@unc.edu (E. Ghysels), jbhill@email.unc.edu (J.B. Hill), motegi@aoni.waseda.jp (K. Motegi).

[1] Early contributions by Zellner and Montmarquette (1971) and Amemiya and Wu (1972) pointed out the potentially adverse effects of temporal aggregation on testing for Granger causality. The subject has been extensively researched ever since, see notably Lütkepohl (1993), Renault et al. (1998), Marcellino (1999), Breitung and Swanson (2002), McCrorie and Chambers (2006), and Silvestrini and Veredas (2008), among others.

http://dx.doi.org/10.1016/j.jeconom.2015.07.007
0304-4076/

to inflation is detected by the MF approach but not when applying conventional Granger causality tests based on LF data. The result suggests that the quarterly frequency is too coarse to capture such causality.

The nature of MF implies that we are potentially dealing with multi-horizon Granger causality since more than one period high frequency (HF) observations are collected within a single LF time span. Moreover, as in the standard (single frequency) VAR literature, exploring MF Granger causality among more than two series also invariably relates to the notion of multi-horizon causality, see in particular Lütkepohl (1993), Dufour and Renault (1998) and Hill (2007). Of direct interest to us is Dufour and Renault (1998) who generalized the original definition of single-horizon or short run causality to multiple-horizon or long run causality to handle causality chains: in the presence of an auxiliary variable, say Z, Y may be useful for a multiple-step ahead prediction of X even if it is useless for the one-step ahead prediction. Dufour and Renault (1998) formalize the relationship between VAR coefficients and multiple-horizon causality and Dufour et al. (2006) formulate Wald tests for multiple-horizon non-causality. Their framework will be used extensively in our analysis.

In addition to the causality literature, the present paper also draws upon and contributes to the MIDAS literature originated by Ghysels et al. (2004, 2005). A number of papers have linked MIDAS regressions to (latent) high frequency VAR models, such as Kuzin et al. (2011) and Foroni et al. (2015), whereas Ghysels (forthcoming) discusses the link between MF-VAR models and MIDAS regressions. None of these papers study in any detail the issue of Granger causality, which is the topic of the present paper.

The paper is organized as follows. In Section 2 we frame MF-VAR models and present core assumptions. In Section 3 we derive the asymptotic properties of the least squares estimator, and a Wald statistic for testing arbitrary linear restrictions on an h-step ahead autoregression. We then develop the mixed frequency causality tests by extending ideas on h-step ahead non-causality tests in Dufour and Renault (1998) and Dufour et al. (2006) to mixed frequencies. Section 4 discusses how we can recover underlying causality using a mixed frequency approach compared to a traditional LF approach. Section 5 shows that the mixed frequency causality tests have higher local asymptotic power than the LF ones do. Section 6 reports Monte Carlo simulation results and documents the finite sample power improvements achieved by the mixed frequency causality test. In Section 7 we apply the mixed frequency and LF causality tests to U.S. macroeconomic data. Finally, Section 8 provides some concluding remarks.

We will use the following notational conventions throughout. Let $\boldsymbol{A} \in \mathbb{R}^{n\times l}$. The l_2-norm is $|\boldsymbol{A}| := (\sum_{i=1}^{n}\sum_{j=1}^{l} a_{ij}^2)^{1/2} = (\text{tr}[\boldsymbol{A}'\boldsymbol{A}])^{1/2}$; the L_r-norm is $\|\boldsymbol{A}\|_r := (\sum_{i=1}^{n}\sum_{j=1}^{l} E|a_{ij}|^r)^{1/r}$; the determinant is $\det(\boldsymbol{A})$; and the transpose is $\boldsymbol{A}'$. $\boldsymbol{0}_{n\times l}$ is an $n\times l$ matrix of zeros. $\boldsymbol{I}_K$ is the K-dimensional identity matrix. Var$[\boldsymbol{A}]$ is the variance–covariance matrix of a stochastic matrix $\boldsymbol{A}$. $\boldsymbol{B}\circ\boldsymbol{C}$ denotes element-by-element multiplication for conformable vectors $\boldsymbol{B}$, $\boldsymbol{C}$.

2. Mixed frequency data model specifications

In this section we present the MF-VAR model and three main assumptions. We want to characterize three settings, respectively high, mixed and low frequency or HF, MF and LF. We begin by considering a partially latent underlying HF process. Using the notation of Ghysels (forthcoming), the HF process contains $\{\{\boldsymbol{x}_H(\tau_L,k)\}_{k=1}^{m}\}_{\tau_L}$ and $\{\{\boldsymbol{x}_L(\tau_L,k)\}_{k=1}^{m}\}_{\tau_L}$, where $\tau_L \in \{0,\ldots,T_L\}$ is the LF time index (e.g. quarterly), $k \in \{1,\ldots,m\}$ denotes the HF (e.g. monthly), and m is the number of HF time periods between LF time indices. In the month versus quarter case, for example, m equals three since one quarter has three months. Observations $\boldsymbol{x}_H(\tau_L,k) \in \mathbb{R}^{K_H\times 1}$, $K_H \geq 1$, are called HF variables, whereas $\boldsymbol{x}_L(\tau_L,k) \in \mathbb{R}^{K_L\times 1}$, $K_L \geq 1$, are latent LF variables because they are not observed at high frequencies—as only some temporal aggregates, denoted $\boldsymbol{x}_H(\tau_L)$, are available.

Note that two simplifying assumptions have implicitly been made. First, there are assumed to be only two sampling frequencies. Second, it is assumed that m is fixed and does not depend on τ_L. Both assumptions can be relaxed at the cost of much more complex notation and algebra which we avoid for expositional purpose—again see Ghysels (forthcoming). In reality the econometrician's choice is limited to MF and LF cases. Only LF variables have been aggregated from a latent HF process in a MF setting, whereas both low *and* high frequency variables are aggregated from the latent HF process to form a LF process. Following Lütkepohl (1987) we consider only linear aggregation schemes involving weights $\boldsymbol{w} = [w_1,\ldots,w_m]'$ such that:

$$\boldsymbol{x}_H(\tau_L) = \sum_{k=1}^{m} w_k \boldsymbol{x}_H(\tau_L,k) \quad\text{and}\quad \boldsymbol{x}_L(\tau_L) = \sum_{k=1}^{m} w_k \boldsymbol{x}_L(\tau_L,k). \tag{2.1}$$

Two cases are of special interest given their broad use: (1) *stock* or *skipped* sampling, where $w_k = I(k=m)$; and (2) *flow* sampling, where $w_k = 1$ for $k=1,\ldots,m$.[2] In summary, we observe:

- all high and low frequency variables $\{\{\boldsymbol{x}_H(\tau_L,j)\}_{j=1}^{m}\}_{\tau_L}$ and $\{\{\boldsymbol{x}_L(\tau_L,j)\}_{j=1}^{m}\}_{\tau_L}$ in a HF process;
- all high frequency variables $\{\{\boldsymbol{x}_H(\tau_L,j)\}_{j=1}^{m}\}_{\tau_L}$ but only aggregated low frequency variables $\{\boldsymbol{x}_L(\tau_L)\}_{\tau_L}$ in a MF process;
- only aggregated high and low frequency variables $\{\boldsymbol{x}_H(\tau_L)\}_{\tau_L}$ and $\{\boldsymbol{x}_L(\tau_L)\}_{\tau_L}$ in a LF process.

A key idea of MF-VAR models is to stack everything observable given a MF process in what we call *the mixed frequency vector*:

$$\boldsymbol{X}(\tau_L) = [\boldsymbol{x}_H(\tau_L,1)',\ldots,\boldsymbol{x}_H(\tau_L,m)',\boldsymbol{x}_L(\tau_L)']'. \tag{2.2}$$

The dimension of the mixed frequency vector is $K = K_L + mK_H$. Note that $\boldsymbol{x}_L(\tau_L)$ is the last block in the mixed frequency vector—a conventional assumption implying that it is observed after $\boldsymbol{x}_H(\tau_L,m)$. Any other order is conceptually the same, except that it implies a different timing of information about the respective processes. We will work with the specification appearing in (2.2) as it is most convenient.

Example 1 (*Quarterly Real GDP*)**.** A leading example of how a mixed frequency model is useful in macroeconomics concerns quarterly real GDP growth $x_L(\tau_L)$, where existing studies of causal patterns use monthly unemployment, oil prices, inflation, interest rates, etc. aggregated into quarters (see e.g., Hill (2007) for references). Consider the monthly oil price changes and CPI inflation stacked into a 6×1 vector (since we have two series for three months) $[\boldsymbol{x}_H(\tau_L,1)',\ldots,\boldsymbol{x}_H(\tau_L,3)']'$, which concatenated with quarterly GDP yields the vector $\boldsymbol{X}(\tau_L)$ appearing in (2.2), which will be further analyzed in Section 7.

We will make a number of standard regulatory assumptions. Let $\mathcal{F}_{\tau_L} \equiv \sigma(\boldsymbol{X}(t) : t \leq \tau_L)$. In particular we assume $E[\boldsymbol{X}(\tau_L)|\mathcal{F}_{\tau_L-1}]$ has a version that is *almost surely* linear in $\{\boldsymbol{X}(\tau_L-1),\ldots,\boldsymbol{X}(\tau_L-p)\}$ for some finite $p\geq 1$.

Assumption 2.1. The process $\boldsymbol{X}(\tau_L)$ is governed by a VAR(p) for some $p\geq 1$:

$$\boldsymbol{X}(\tau_L) = \sum_{k=1}^{p} \boldsymbol{A}_k \boldsymbol{X}(\tau_L-k) + \boldsymbol{\epsilon}(\tau_L). \tag{2.3}$$

[2] One can equivalently let $w_k = 1/m$ for $k=1,\ldots,m$ in flow sampling if the average is preferred to a summation.