ARTICLE IN PRESS

International Journal of Forecasting ▮ (▮▮▮▮) ▮▮▮–▮▮▮

Contents lists available at ScienceDirect

International Journal of Forecasting

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

Real-time nowcasting the US output gap: Singular spectrum analysis at work

Miguel de Carvalho[a], António Rua[b,c,*]

[a] Faculty of Mathematics, Pontificia Universidad Católica de Chile, Santiago, Chile

[b] Economics and Research Department, Banco de Portugal, Lisboa, Portugal

[c] Nova School of Business and Economics, Universidade Nova de Lisboa, Lisboa, Portugal

ARTICLE INFO

Keywords:
Band-pass filter
Fisher *g* test
Multivariate singular spectrum analysis
Singular spectrum analysis
US output gap
Real-time data

ABSTRACT

We explore a new approach for nowcasting the output gap based on singular spectrum analysis. Resorting to real-time vintages, a recursive exercise is conducted in order to assess the real-time reliability of our approach for nowcasting the US output gap, relative to some well-known benchmark models. For our application of interest, the preferred version of our approach is a multivariate singular spectrum analysis, where we use a Fisher *g* test to infer which components, within the standard business cycle range, should be included in the grouping step. We find that singular spectrum analysis provides a reliable assessment of the cyclical position of the economy in real time, with the multivariate approach outperforming its univariate counterpart substantially.

1. Introduction

The output gap plays a central role in policymaking. Most central banks aim to keep inflation under control, and the output gap is a key source of inflation pressure in the economy. Given that the output gap fluctuates when the economy is overheating or underperforming, the conduct of monetary policy should take it into full consideration. Since the cyclical position of the economy may influence fiscal policy, it can also be used by governments to help determine and pursue policy measures; thus, the assessment of the output gap is crucial in many cases for the formulation of countercyclical stabilization policy.

However, measuring the output gap is challenging, as it cannot be observed directly, and therefore cannot be assessed precisely. The revisions to which real-time output gap estimates are subject form yet another challenge, as they can compromise the operational usefulness of the estimates for policymakers, who need reliable 'intel' in real-time. To date, there have been several studies that have documented the large uncertainty of real-time output gap estimates, with this being a common issue for all estimation methods available (see Edge & Rudd, 2012; Marcellino & Musso, 2011; Orphanides, 2003a; Orphanides & van Norden, 2002; Watson, 2007, among others). The policy implications of the effects of output gap uncertainty have been addressed by, for example, Orphanides (2001, 2003b), Orphanides and Williams (2007), Rudebusch (2001), and Smets (2002).

In this paper, we focus on singular spectrum analysis (SSA), and evaluate its potential contribution for nowcasting the output gap in a real-time setup. Despite the potential usefulness of SSA for the analysis of economic phenomena, there have been only a few applications in the economics and finance literature. In this respect, see the recent work by de Carvalho, Rodrigues, and Rua (2012),

* Correspondence to: Economics and Research Department, Banco de Portugal, Av. Almirante Reis no. 71, 1150–012 Lisboa, Portugal.

E-mail addresses: MdeCarvalho@mat.puc.cl (M. de Carvalho), Antonio.Rua@bportugal.pt (A. Rua).

http://dx.doi.org/10.1016/j.ijforecast.2015.09.004
0169-2070/

Please cite this article in press as: de Carvalho, M., & Rua, A., Real-time nowcasting the US output gap: Singular spectrum analysis at work. International Journal of Forecasting (2016), http://dx.doi.org/10.1016/j.ijforecast.2015.09.004

Hassani, Heravi, and Zhigljavsky (2009), Hassani, Heravi, and Zhigljavsky (2013), Hassani, Soofi, and Zhigljavsky (2013), and Patterson, Hassani, Heravi, and Zhigljavsky (2011).

The implementation of SSA involves the selection of two important parameters. In the decomposition stage, one has to set the window length, L, in the embedding step, whereas in the reconstruction stage, one has to choose the number of components for the grouping step. Concerning the choice of L, as pointed out by Hassani, Mahmoudvand, and Zokaei (2011), large values of L allow longer period oscillations to be resolved, but if L is too large, it leaves too few observations from which to estimate the covariance matrix of the L variables. Golyandina, Nekrutkin, and Zhigljavsky (2001) recommended that L should be large enough but no larger than $T/2$, whereas Elsner and Tsonis (1996) discuss the practice of choosing L equal to $T/4$, where T is the sample size. Hassani et al. (2009) argue that if the time series presents a periodic component with an integer period, it is advisable to take a window length that is proportional to that period, in order to achieve a better separability of the periodic component. Drawing on the concept of separability between the signal and noise components, Hassani et al. (2011) find that a suitable value for L, at least for reconstruction purposes, is close to Median$\{1, \ldots, T\}$ for a series of length T (see also Hassani, Mahmoudvand, Zokaei, & Ghodsi, 2012). However, as mentioned by Hassani, Heravi, and Zhigljavsky (2013), such a value may not be optimal for forecasting purposes (Mahmoudvand, Najari, & Zokaei, 2013). In this respect, Hassani, Soofi et al. (2013) take into account the forecasting horizon of interest when selecting L, whereas Hassani, Webster, Silva, and Heravi (2015) choose the value of L that minimizes the forecasting error. In the context of output gap estimation, de Carvalho et al. (2012) suggest selecting L as the maximum period of the business cycle frequency range that one is interested in. As stressed by Hassani (2007), the selection of the proper window length depends upon the problem at hand, and on preliminary information about the time series. Finally, it should be noted that rules that may be optimal for (univariate) SSA need not be optimal for multivariate SSA (MSSA). For example, the above-mentioned rule based on Median$\{1, \ldots, T\}$ is suboptimal for MSSA; cf. Hassani and Mahmoudvand (2013, Proposition 1).

The selection of the components to be used in the reconstruction stage is also far from straightforward. In the separation of the signal and noise, one way to proceed is by looking at the plot of the eigenvalues, properly ordered by values, associated with the reconstructed components. This plot can hopefully guide the truncation of the number of components to be considered in the grouping step. As was mentioned by Hassani et al. (2009), a slowly decreasing sequence of eigenvalues is usually related to noise, whereas similar values of the eigenvalues allow the identification of the eigentriples that correspond to the same harmonic component of the series. Furthermore, they suggest computing the periodogram to assist in selecting the components to group. In fact, the presence of peaks in the periodogram provides an indication of the harmonic components in the series. Another possible approach is to compute w-correlations between the components; see Golyandina et al. (2001) and Hassani (2007). Low values for w-correlations between the reconstructed components indicate that the components are well separated, whereas high values suggest that they should be considered as a group and may pertain to the same component in the SSA decomposition. Among other alternative approaches, one should mention that of Hassani et al. (2015), who choose the number of components that minimizes the forecasting error (see also Cassiano et al., 2013).

Recently, de Carvalho et al. (2012) have shown that SSA can deliver output gap estimates that resemble those obtained using band-pass filters, while improving the reliability of the corresponding nowcasts. Here, we extend their work in several dimensions. First, to mimic a real-life policymaking scenario, that is, to replicate the problems faced by policymakers at the time when policy decisions have to be taken, we consider real-time data. This means considering the vintages of data that were available at each moment in time. It is widely acknowledged by now that data revisions can affect policy decisions, and although the issue of the importance of data revisions is not recent, there has been a growing level of interest among practitioners in the inclusion of real-time data in analyses, since the influential work by Croushore and Stark (2001, 2003), who compiled and examined real-time data for major US macroeconomic variables. Hence, we focus on the evaluation of output gap nowcasts computed through a recursive exercise using the vintage of data available at each period. This allows us to obtain real-time estimates, which are the ones that are relevant in terms of policymaking, whereas de Carvalho et al. (2012) only considered *quasi*-real estimates, by considering the latest available vintage.[1]

Second, we suggest a novel approach for the selection of the principal components to be used for reconstructing the cyclical component of a variable of interest. de Carvalho et al. (2012) use a heuristic approach to select the components to be considered for the reconstruction of the cyclical component of GDP. Based on the dominant frequency, they consider the components that reflect periodicities of interest, namely those within the business cycle frequency range. In this respect, Hassani et al. (2009) suggest the computation of the periodogram for assessing the dominant periodicity. We address this issue through the proposal of an alternative inferential procedure that uses a spectral-based Fisher g test. Although Fourier analysis is less popular than time domain analysis, it has proven to be quite useful in a wealth of contexts (see, for example, A'Hearn & Woitek, 2001; Breitung & Candelon, 2006; Lemmens, Croux, & Dekimpe, 2008; Rua & Nunes, 2005). Drawing on the periodogram estimator, Fisher (1929) derived an exact test – the so-called Fisher g test – which allows for the detection of hidden periodicities of an unspecified frequency, by determining whether a peak in the periodogram is significant or not. We use the Fisher g test

[1] The *quasi*-real estimate is the rolling estimate based on the latest available vintage. Real-time and *quasi*-real estimates both cover the same period, and differ only due to data revisions. See Orphanides and van Norden (2002, p. 571).