International Economics 166 (2021) 42–57

Contents lists available at ScienceDirect

International Economics

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

Frequency domain analysis and filtering of business and consumer survey expectations

Oscar Claveria [a,*], Enric Monte [b], Salvador Torra [c]

[a] AQR-IREA (Institute of Applied Economics Research), University of Barcelona (UB), Diagonal, 690, 08034, Barcelona, Spain
[b] Department of Signal Theory and Communications, Polytechnic University of Catalunya (UPC), Jordi Girona, 1-3, 08034, Barcelona, Spain
[c] Riskcenter-IREA, Department of Econometrics, Statistics and Applied Economics, University of Barcelona, Diagonal, 690, 08034, Barcelona, Spain

ARTICLE INFO

JEL classification:
C65
C82

Keywords:
Business and consumer surveys
Spectral analysis
Seasonality
Signal processing
Low-pass filter

ABSTRACT

The main objective of this study is two-fold. First, we aim to detect the underlying existing periodicities in business and consumer survey expectations through spectral analysis. We use the Welch method to extract the harmonic components that correspond to the cyclic and seasonal patterns in all response options of monthly survey indicators. We find that business survey indicators show a common cyclical component of low frequency that corresponds to about four years, while for most consumer survey indicators, we do not detect any relevant cyclic components, and the obtained lower frequency periodicities show a very irregular pattern across questions and reply options. Second, we aim to provide researchers with a filter specially designed for business and consumer survey expectations that circumvents other filtering methods' assumptions. Most seasonal adjustment methods are based on *a priori* assumptions about the structure of the components and do not depend on the features of the specific series. To overcome this limitation, we design a low-pass filter that allows extracting the components with periodicities similar to those that can be found in the dynamics of economic activity. We opt for a Butterworth filter and apply a zero-phase filtering process to preserve the time series' time alignment. We find that the balances computed with the filtered series are highly correlated with the seasonally-adjusted balances published by the European Commission, albeit the former tend to be smoother for the consumer survey indicators.

1. Introduction

Sentiment analysis is increasingly used for anticipating the evolution of the economy and for assisting in the design of economic policies and business decision making. One of the main sources of information for sentiment analysis is agents' expectations about future developments of the economy. Every month, the European Commission (EC) publishes the European Sentiment Indicator (ESI), which synthesises the beliefs and expectations of economic agents. The ESI is a composite indicator based on business and consumer survey results, so it offers a perspective both from the supply side and from the demand side of the economy.

In the surveys, respondents are asked about their expectations regarding a wide range of economic variables, and they are mostly faced with three reply options: "up", "unchanged" and "down". Results are published as balances, which can be regarded as diffusion indexes consisting of the subtraction between the aggregate percentages of response corresponding to the extreme categories.

* Corresponding author. Department of Econometrics, Statistics and Applied Economics, University of Barcelona, Diagonal 690, 08034, Barcelona, Spain.

E-mail addresses: oclaveria@ub.edu (O. Claveria), enric.monte@upc.edu (E. Monte), storra@ub.edu (S. Torra).

https://doi.org/10.1016/j.inteco.2021.03.002
Received 18 September 2020; Received in revised form 26 February 2021; Accepted 3 March 2021
Available online 20 March 2021

Seasonally-adjusted balances are published each month by the EC, but the series corresponding to each response category are only available in raw form, that is, the aggregate percentage of respondents in each category.

The main objective of this study is two-fold. On the one hand, we aim to detect the underlying existing periodicities in business and consumer survey expectations. With this objective, we conduct a spectral analysis of all monthly survey information from the EC industry and consumer surveys in the euro area (EA). Given that the periodic components may vary across survey questions, this approach gives insight into the specificities of the different survey indicators and also provides researchers with information to improve the design of ad-hoc filters for signal processing of survey data (Gottman, 1981).

On the other hand, we aim to provide researchers with a filter for business and consumer survey data that allows circumventing the application of filters based on *a priori* assumptions on the structure of the components. For this purpose, we design a low-pass filter that allows extracting the components with periodicities similar to those that can be found in the dynamics of economic activity.

In this research we apply the Welch method (Welch, 1967; Oppenheim and Schafer, 2010) for spectral density estimation to detect periodicities in the time domain. Welch's algorithm can be regarded as spectral density estimation procedure that allows reducing the high variability in the estimation obtained by means of the periodogram. This reduction of noise in the estimated power spectra is particularly desirable when dealing with finite data, as is the case for survey indicators.

The fact that most methods for seasonal adjustment are based on assumptions about the structure of the components and, do not depend on the features of the series, has led us to design a low-pass filter for survey indicators. We use a Butterworth filter and apply a zero-phase filtering process to preserve the time alignment of the series. This procedure allows us to reject the frequency components of the survey indicators that do not have a counterpart in the dynamics of economic activity. Finally, we assess the performance of the filtered survey indicators by comparing them to the seasonally-adjusted balances published by the EC.

The rest of the paper proceeds as follows. Section 2 describes the data. Section 3 presents the methodological approach and provides the results of the spectral analysis of survey indicators. In Section 4, we present a filter for business and consumer survey data and describe its design and properties. Finally, concluding remarks are given in Section 5.

2. Data

This study uses qualitative expectations of firms and consumers from within the EA about a wide array of economic variables. Specifically, we use all monthly raw data from the joint harmonised EU industry and consumer surveys conducted by the EC (see Table 1) (European Commission, 2019). The sample period goes from 1993.M1 to 2019.M2. In the industry survey, manufacturers are asked about their expectations regarding firm-specific factors such as production, selling prices and employment and, they are faced with three options: "up", "unchanged" and "down". P_t measures the percentage of respondents reporting an increase in the variable, E_t no change, and M_t a decrease. The most common way of presenting survey data is the balance, B_t, which is computed as the subtraction between the two extreme categories:

$$B_t = P_t - M_t \tag{1}$$

Consumers, for their part, are asked about objective variables (e.g. how they think the general economic situation, the cost of living, and the level of unemployment in the country will change over the next twelve months) and subjective variables (e.g. major purchases,

Table 1
Survey indicators.

Industry survey
Monthly questions
INDU.1 – Production trend observed in recent months
INDU.2 – Assessment of order-book levels
INDU.3 – Assessment of export order-book levels
INDU.4 – Assessment of stocks of finished products
INDU.5 – Production expectations for the months ahead
INDU.6 – Selling price expectations for the months ahead
INDU.7 – Employment expectations for the months ahead
Consumer survey
Monthly questions
CONS.1 – Financial situation over last 12 months
CONS.2 – Financial situation over next 12 months
CONS.3 – General economic situation over last 12 months
CONS.4 – General economic situation over next 12 months
CONS.5 – Price trends over last 12 months
CONS.6 – Price trends over next 12 months
CONS.7 – Unemployment expectations over next 12 months
CONS.8 – Major purchases at present
CONS.9 – Major purchases over next 12 months
CONS.10 – Savings over next 12 months
CONS.11 – Statement on financial situation of household

savings, etc.). Consumers have three additional response categories: two at each end ("a lot better/much higher/sharp increase", and "a lot worse/much lower/sharp decrease"), and a "don't know" option. As a result, PP_t measures the percentage of respondents reporting a sharp increase in the variable, P_t a slight increase, E_t no change, M_t a slight fall, MM_t a sharp fall and, N_t don't know. See Gelper and Croux (2010) for an appraisal of the data from the EU business and consumer surveys. Again, the most common way of presenting survey data is the balance, which in the case of five reply options is computed as a weighted mean as follows:

$$B_t = \left(PP_t + {}^1\!/\!{}_2P_t\right) - \left({}^1\!/\!{}_2M_t + MM_t\right) \tag{2}$$

Business and consumer survey data record opinions of agents, and as such, they are influenced by unpredictable economic and political events. As stated by the EC, even when respondents are required to indicate whether the reported change can be attributed to seasonal factors, time series often display characteristic seasonal fluctuations (European Commission, 2016). As a result, the EC applies a specific method for seasonal adjustment known as Dainties, which was originally developed by Eurostat. Dainties uses filters that are based on the assumption of the breakdown of the series in three components (seasonal, trend and irregular), and applies filters that are decided ex-ante and that do not depend on the data.

In this paper we propose a frequency domain approach that avoids making any *a priori* assumption regarding the structure of the components. We analyse the data directly in order to detect the different periodic components. The proposed approach uses decomposition in a sum of periodic functions such as sines and cosines. These components are estimated by means of a variation of the Fourier transform, where the trend component is reflected in the extreme of low frequency components.

3. Spectral analysis

In this section we describe the methodological approach used for transforming the information from the time domain to the frequency domain and present the results of the spectral analysis for the detection of periodicities in survey indicators. The proposed procedure is based on the Welch method for spectral density estimation. This approach allows us to gain some insight into potential hidden structures in business and consumer survey data.

Discrete Fourier Transform (DFT) can be regarded as a mathematical tool that when applied to any discrete time series decomposes them into a linear combination of discrete frequencies. These frequencies determine the importance of a given cyclical component of the time series. Thus, a frequency component with a higher value in the DFT will correspond to a cyclical component with higher energy (amplitude). The statistical average of a certain signal as analysed in terms of its frequency content is called its spectrum. Fourier analysis allows converting a signal from its original domain, time in our case, to a representation in the frequency domain, and vice versa.

In signal processing, a periodogram is an estimate of the spectral density of a signal, and is computed as the square of the absolute value of the Fourier transform of the time series. The periodogram is a standard method for determining the spectral components of a time series, and it is the most common tool for examining the amplitude vs frequency characteristics. Periodograms have been used for more than one century (Gardner, 1987). Nevertheless, one drawback of the periodogram is that it presents a very high variability in terms of the estimation at each frequency.

Welch's method represents an improvement on the standard periodogram spectrum estimation in that it reduces noise in the estimated power spectra in exchange for reducing the frequency resolution. The method is based upon the usage of periodogram spectrum estimates, which are the result of converting a signal from the time domain to the frequency domain. Although similar to Bartlett's method (Bartlett, 1948, 1950), it differs in that: i) the signal is split up into overlapping data segments, ii) the overlapping segments are then windowed in the time domain, resulting in a modified periodogram. As a result, a high resolution in frequency with high amplitude noise is traded for a lower resolution in frequency with reduced amplitude noise, diminishing the estimation variability.

While most windowing functions assign higher weight to the data at the centre of the set than to the data at the edges, which implies a loss of information, Welch's method mitigates that loss by overlapping in time. In this sense, we also want to note the potential of new signal processing techniques such as Prism (precise, repeat integral signal monitor), which is immune to frequency discretization. See Henry (2021) for a detailed description.

This procedure is done as follows. First, the original time series are fragmented into overlapping segments. Then, for each segment a periodogram is computed by means of DFT of business and consumer survey indicators, denoted as x_t, where $x_t = \{x_1,\ldots,x_T\}$. The DFT transforms a sequence x_t into another sequence of complex numbers $X_k = \{X_1,\ldots,X_T\}$ by means of the following expression:

$$F(x_t) = \sum_{t=1}^{T} x_t e^{-\frac{i2\pi}{T}kt} = \sum_{t=1}^{T} x_t \left[\cos\left(\frac{2\pi}{T}kt\right) - i\cdot\sin\left(\frac{2\pi}{T}kt\right)\right] \tag{3}$$

where $F(x_t)$ denotes the Fourier component at frequency k; i denotes the imaginary number square root of -1; t is a time index of the sequence $t = 1,\ldots,T$; and the frequency component k indicates the inverse of the detected periodicity, i.e. T/k. Finally, the squared magnitude of the result is computed and the individual periodograms are averaged using the Welch method to reduce the variance of the individual power measurements.

Spectral density analysis tools developed in the field of signal processing do not provide confidence margins (Oppenheim and Schafer, 2010). In order to discern the presence of a periodic component with respect to the background noise, the decision is based on the component's energy and its physical viability, and it is circumscribed to the type of analysis.The energy of a component can be

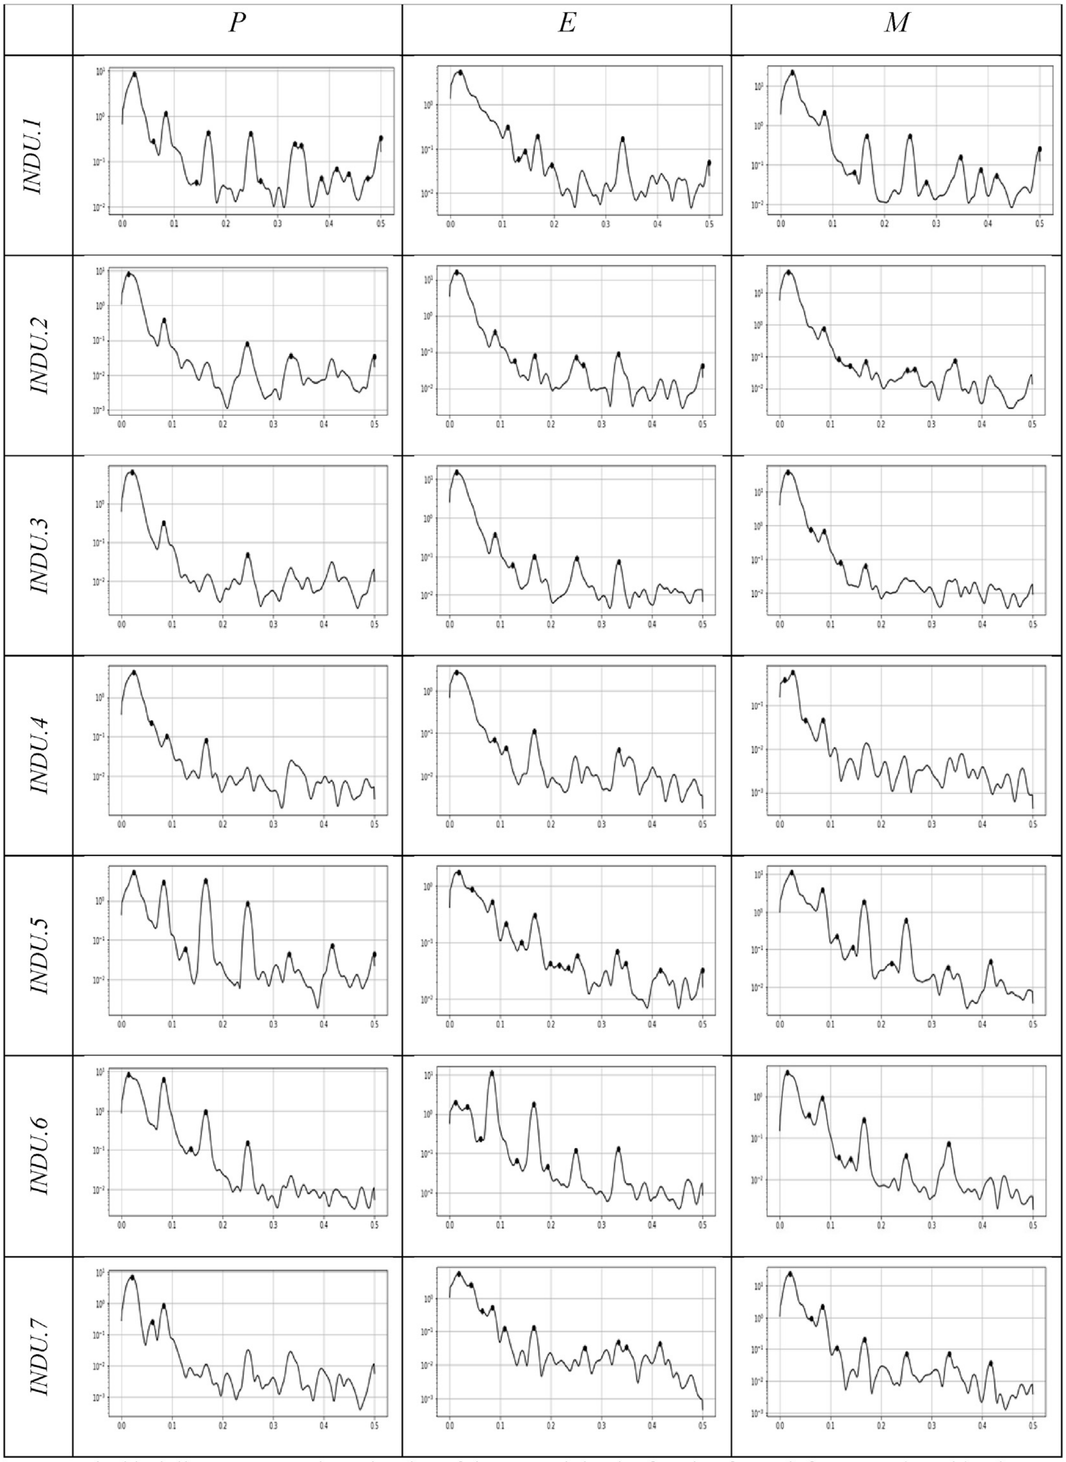

Notes: The black line represents the estimation of the spectral density function for each frequency (Y-axis). The X-axis the normalised frequencies, measured in cycles per month. The dots indicate the cycles detected in the reference time series using the built-in method for automatic peak detection that comes with Python's scipy library.

Fig. 1. Spectral density functions of the industry survey indicators – Monthly questions. Notes: The black line represents the estimation of the spectral density function for each frequency (Y-axis). The X-axis the normalised frequencies, measured in cycles per month. The dots indicate the cycles detected in the reference time series using the built-in method for automatic peak detection that comes with Python's scipy library.

defined as its relative amplitude with respect to the background noise level. Consequently, the decision on whether to consider the presence of a periodic component as relevant has an interpretative factor.

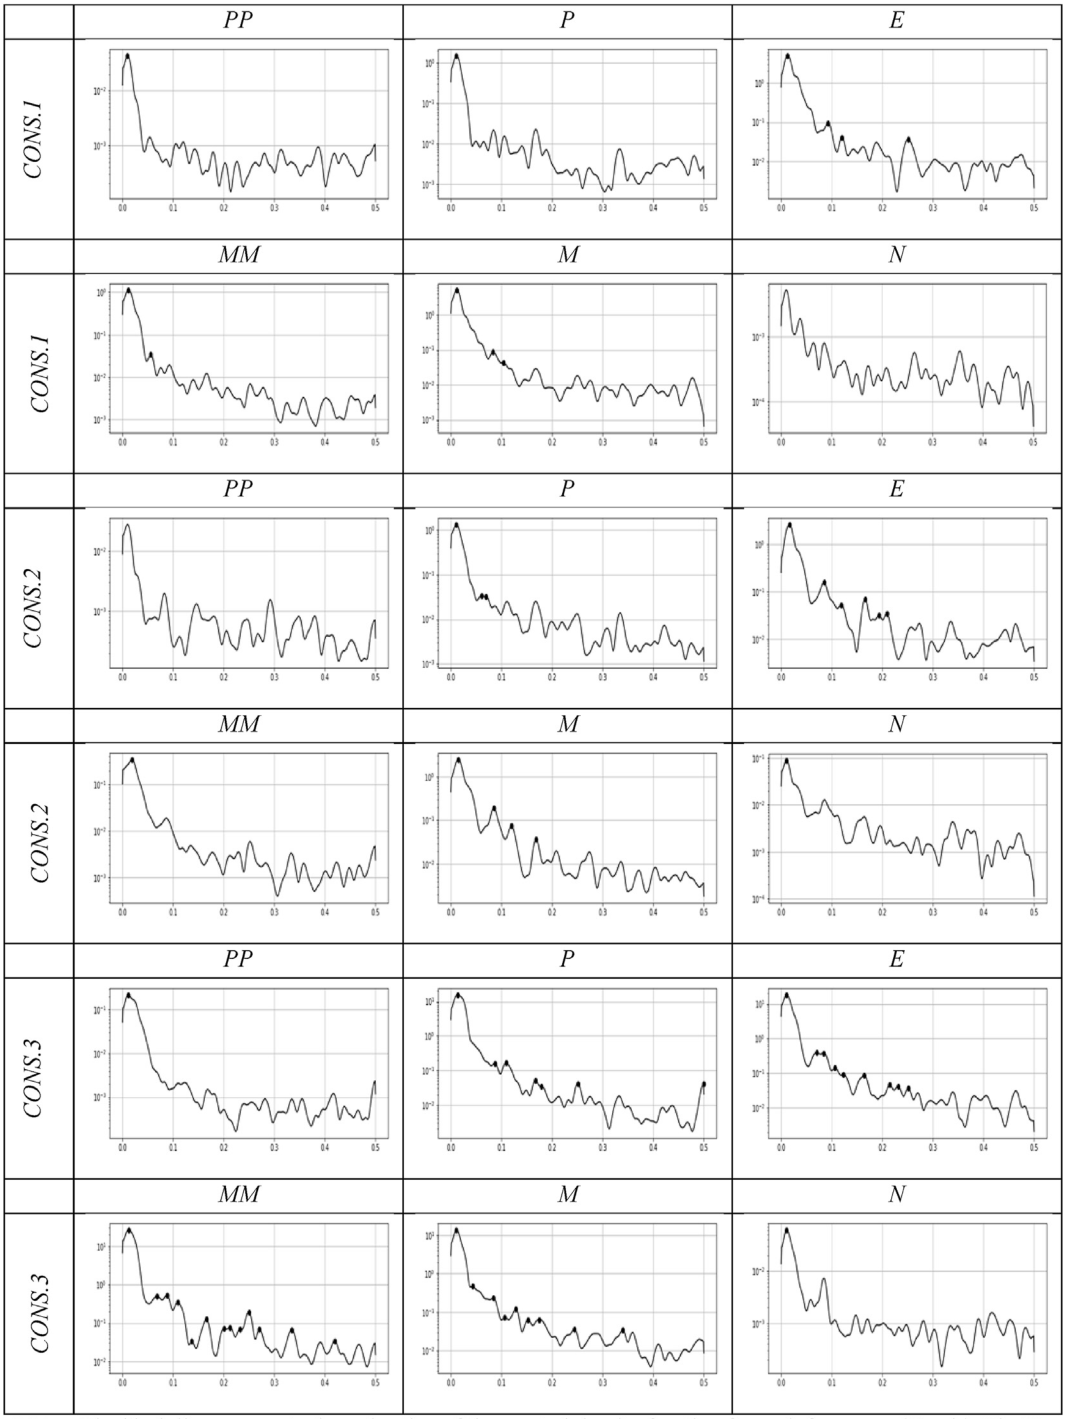

Notes: The black line represents the estimation of the spectral density function for each frequency (Y-axis). The X-axis the normalised frequencies, measured in cycles per month. The dots indicate the cycles detected in the reference time series using the built-in method for automatic peak detection that comes with Python's scipy library.

Fig. 2. Spectral density functions of the consumer survey indicators – Monthly questions. Notes: The black line represents the estimation of the spectral density function for each frequency (Y-axis). The X-axis the normalised frequencies, measured in cycles per month. The dots indicate the cycles detected in the reference time series using the built-in method for automatic peak detection that comes with Python's scipy library.

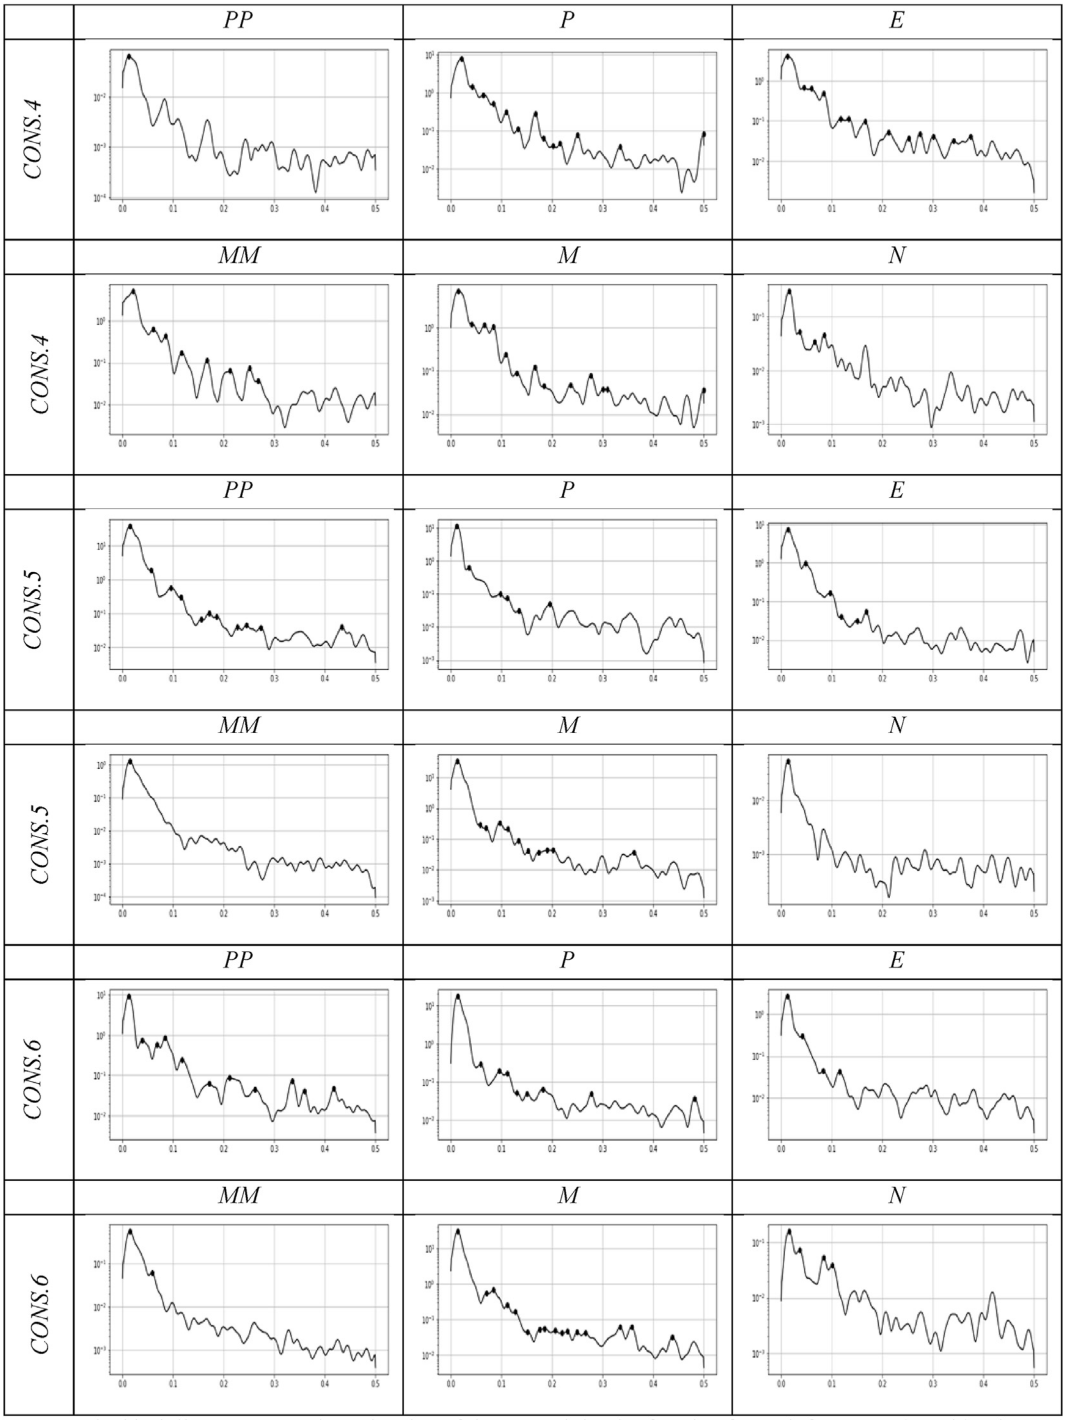

Notes: The black line represents the estimation of the spectral density function for each frequency (Y-axis). The X-axis the normalised frequencies, measured in cycles per month. The dots indicate the cycles detected in the reference time series using the built-in method for automatic peak detection that comes with Python's scipy library.

Fig. 3. Spectral density functions of the consumer survey indicators – Monthly questions. Notes: The black line represents the estimation of the spectral density function for each frequency (Y-axis). The X-axis the normalised frequencies, measured in cycles per month. The dots indicate the cycles detected in the reference time series using the built-in method for automatic peak detection that comes with Python's scipy library.

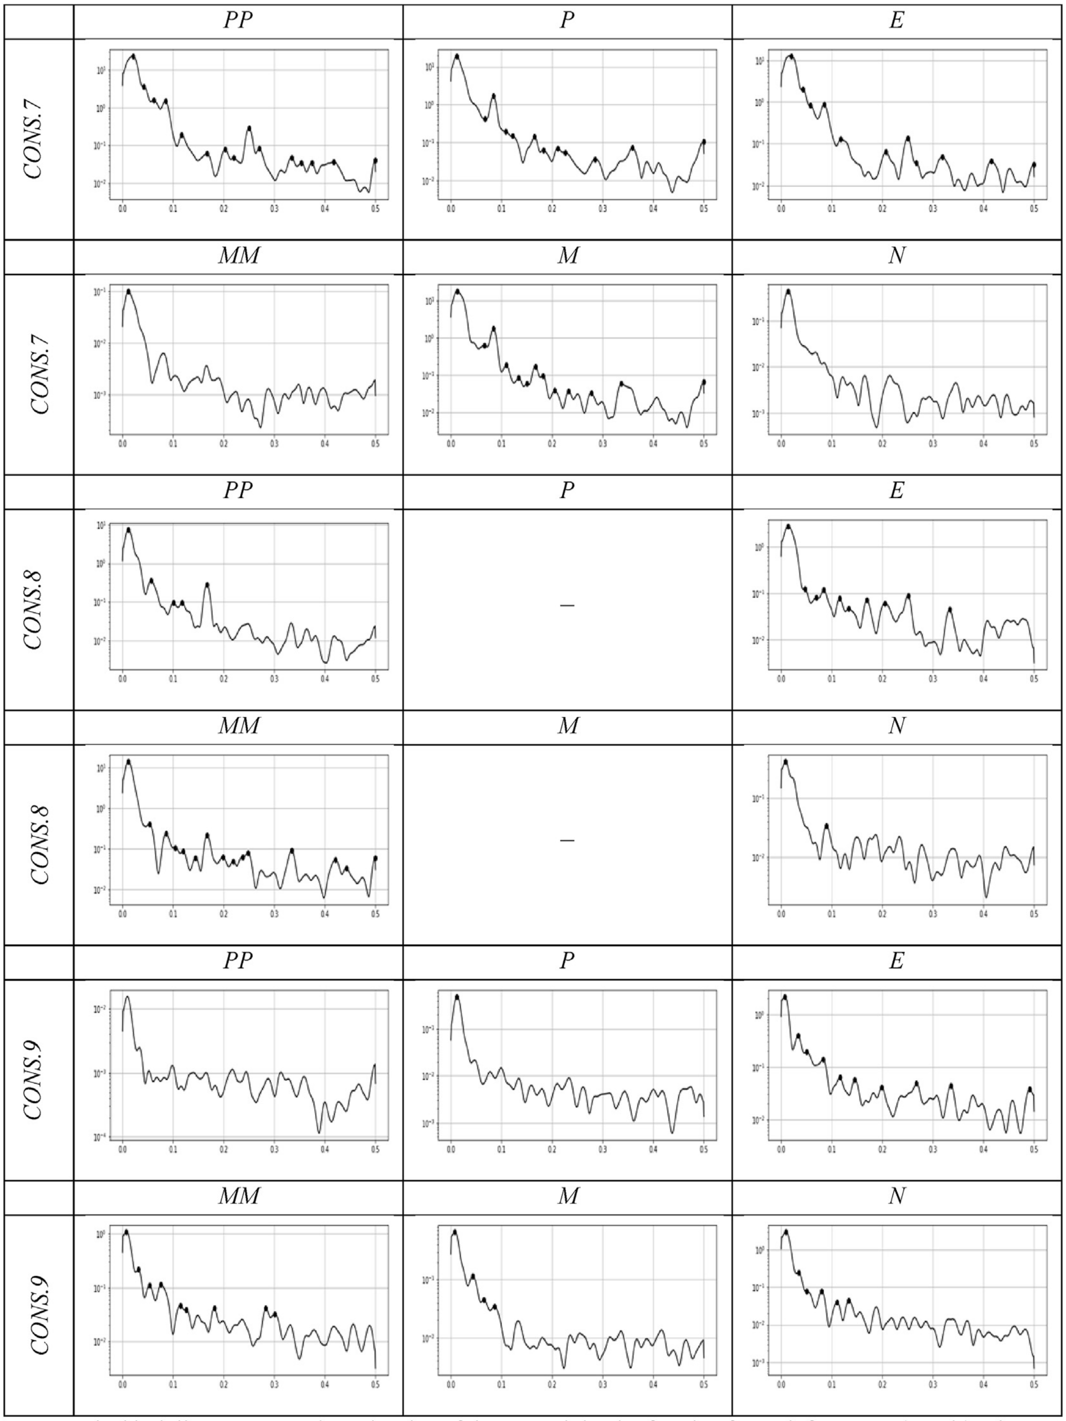

Notes: The black line represents the estimation of the spectral density function for each frequency (Y-axis). The X-axis the normalised frequencies, measured in cycles per month. The dots indicate the cycles detected in the reference time series using the built-in method for automatic peak detection that comes with Python's scipy library. Survey question *CONS.8* only has four reply options; N/A categories (*P* and *M*) are marked with a – sign.

Fig. 4. Spectral density functions of the consumer survey indicators – Monthly questions. Notes: The black line represents the estimation of the spectral density function for each frequency (Y-axis). The X-axis the normalised frequencies, measured in cycles per month. The dots indicate the cycles detected in the reference time series using the built-in method for automatic peak detection that comes with Python's scipy library. Survey question *CONS.8* only has four reply options; N/A categories (*P* and *M*) are marked with a – sign.

We have used the following criterion. Since the length of the series is 324 months, to have an average of 4 windows with 50% overlap, we used a window length (L) of 120 observations. We choose a rectangular window. In this way, on the one hand, we obtain a resolution of $1/L$ for each window, understood as the ability to separate two sinusoidal waves. On the other hand, the reduction in estimation noise is half, as it is proportional to $1/\sqrt{4}$. As a result, the lowest frequency sinusoid (maximum period) that can be clearly detected is 60 months. The criterion for deciding whether the peak is relevant is that the region at a distance greater than $\pm 1/60$ should have a power level of less than half. For intermediate frequency peaks, we follow a similar criterion. On the one hand, the separation between peaks is set to be greater to $1/60$ in discrete frequency and, on the other hand, that between two successive peaks there is a valley with a level less than half of the lowest amplitude peak.

Next, we present the results of applying the Welch algorithm to business and consumer survey data. We have computed the spectral components of time series with the built-in method for automatic peak detection that comes with Python's scipy library (https://scipy.org/). These peaks correspond to the underlying periodic components of each survey indicator. Results are grouped by questions.

In each graph of Figs. 1–5, we present the estimates of the spectral density for each response category of each question. In the first subsection (3.1) we present the results for the industry survey and, in the second subsection (3.2) the results for the consumer survey. The Y axis represents the DFTs (3), which are expressed on a logarithmic scale. The X axis represents the normalised frequencies, measured in cycles per month.

The dots in each graph indicate the periodic components detected in the reference time series. As explained before, the importance of the peaks is given by their relative local variation, giving rise to peaks that cannot be considered relevant. As a result, only some of the

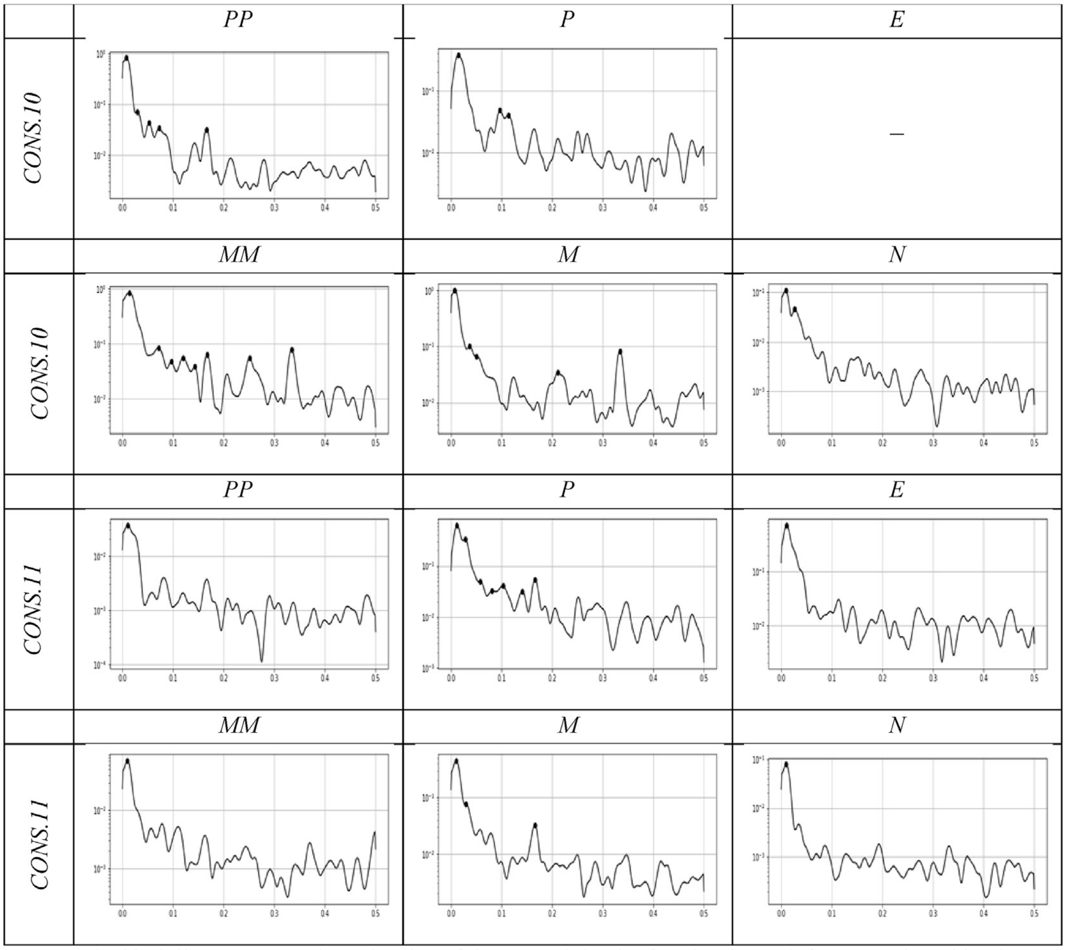

Notes: The black line represents the estimation of the spectral density function for each frequency (Y-axis). The X-axis the normalised frequencies, measured in cycles per month. The dots indicate the cycles detected in the reference time series using the built-in method for automatic peak detection that comes with Python's scipy library. Survey question *CONS.10* only has five reply options; N/A categories (*E*) are marked with a – sign.

Fig. 5. Spectral density functions of the consumer survey indicators – Monthly questions. Notes: The black line represents the estimation of the spectral density function for each frequency (Y-axis). The X-axis the normalised frequencies, measured in cycles per month. The dots indicate the cycles detected in the reference time series using the built-in method for automatic peak detection that comes with Python's scipy library. Survey question *CONS.10* only has five reply options; N/A categories (*E*) are marked with a – sign.

peaks are clearly associated with the underlying periodicities. To facilitate the interpretation of the obtained results, in Tables 2 and 3 we present the relevant periodic components detected in Figs. 1–5, conveniently re-expressed in months. Table 2 summarises the main periodicities for the industry indicators and Table 3 the main ones for the consumer indicators.

3.1. Industry survey

In Table 2 we observe that most survey indicators tend to present a strong cyclical component of about 4 to 4½ years, with the exception of category *E* in the case of 'selling price expectations for the months ahead' (*INDU.6*).

Other common periodicities between all reply options from all questions are the ones corresponding to 11–12 months, half a year, a quarter of a year, and 3 months. These periodic components were clearly detected as high energy peaks in the Fourier transforms, especially for category *P* in two questions: 'production expectations for the months ahead' (*INDU.5*) and 'selling price expectations for the months ahead' (*INDU.6*).

3.2. Consumer survey

Results in Table 3 show that for consumer survey expectations the intensity of the cycles heavily varies across questions and reply options. In most cases, we do not detect any cyclical components, and when we do, these peaks are of very low energy so as to be considered relevant. One of the questions for which we obtain clearly marked peaks is the one regarding the 'financial situation over next 12 months' (*CONS.2*), especially for reply option *M*. Curiously, when the question is stated in the past, regarding the last 12 months (*CONS.2*), the result is opposite: the detected peaks are of very little energy in all categories. This difference between forward-looking questions (expectations) and those referring to the past (perceptions), is repeated in some other cases. For example, prospective questions such as the one referring to the 'general economic situation over the next 12 months' (*CONS.4*) tend to show more relevant lower frequency periodicities than when they are referred to the past (*CONS.3*).

In summary, we find that the detected lower frequency periodicities for consumer survey indicators show a very irregular pattern across questions and reply options. This pattern notably differs from the results obtained for industry survey indicators (Table 2). Business survey indicators showed a common cyclical component of low frequency that corresponds to about 4 to 4½ years which was present across reply options from all questions.

It should be highlighted that to a certain extent these findings may be conditioned by several biases derived from the fact that the wording of the questions differs between both surveys. While business survey indicators mostly refer to the evolution of the variables 'in the recent months' or 'for the months ahead', consumer survey indicators refer to the situation 'over the last twelve months' or 'over the next twelve months'. Furthermore, as opposed to firms, which are faced with three reply options, consumers have three additional response categories: two at each end ("a lot better/much higher/sharp increase", and "a lot worse/much lower/sharp decrease"), and a "don't know" option. Additionally, we want to point out the different nature between the questions of both surveys, since the questions of the consumer survey refer to objective variables, while the questions of the business survey refer to specific factors of the firm. Finally, as suggested by Kahneman and Tversky (1971), there may be additional biases attributable to the agents themselves. Phenomena such as loss aversion or anchoring may be affecting respondents when they answer the questionnaires, and should not be neglected.

Table 2
Periodicities of the industry survey indicators – Monthly questions.

Variable	Reply option	Periodicities (months)
INDU.1 – Production trend observed in recent months	*P*	56.9, 11.1, 6.0, 4.0, 3.0
	E	56.9, 5.9, 3.0
	M	56.9, 6.0, 4.0
INDU.2 – Assessment of order-book levels	*P*	52.9, 12.0, 4.0, 3.0
	E	52.9, 3.0
	M	52.9, 11.4
INDU.3 – Assessment of export order-book levels	*P*	52.9, 11.4, 4.0
	E	52.9, 11.6, 6.0, 4.3
	M	52.9, 11.6, 6.0
INDU.4 – Assessment of stocks of finished products	*P*	52.9, 6.0,
	E	56.9, 6.0, 2.9
	M	50.0, 11.0
INDU.5 – Production expectations for the months ahead	*P*	52.9, 11.4, 6.0, 4.0
	E	52.9, 6.0
	M	52.9,11.4, 6.0, 4.0, 2.9
INDU.6 – Selling price expectations for the months ahead	*P*	52.9, 11.4, 6.0, 4.0
	E	11.4, 6.0, 4.0, 3.0
	M	52.9, 11.4, 6.0, 4.0, 3.0
INDU.7 – Employment expectations for the months ahead	*P*	48.8,12.1
	E	56.9, 12.2, 6.0
	M	52.9, 11.9, 5.9, 4.0, 2.9

Table 3
Periodicities of the consumer survey indicators – Monthly questions.

Variable	Reply option	Periodicities
CONS.1 – Financial situation over last 12 months	*PP*	–
	P	–
	E	10.8, 8.3, 4.0
	M	12.0
	MM	18.0
	N	–
CONS.2 – Financial situation over next 12 months	*PP*	–
	P	–
	E	11.8, 6.0
	M	11.8, 8.3, 6.0
	MM	–
	N	–
CONS.3 – General economic situation over last 12 months	*PP*	–
	P	–
	E	11.9, 8.1, 6.1
	M	–
	MM	11.4, 6.0, 4.3
	N	12.0
CONS.4 – General economic situation over next 12 months	*PP*	12.0
	P	48.7, 6.0, 3.0
	E	–
	M	68.3, 12.0, 9.2, 6.0, 4.2, 3.6
	MM	46.6, 16.5, 11.8, 8.5, 6.0, 4.7, 4.0
	N	64.0, 11.8
CONS.5 – Price trends over last 12 months	*PP*	–
	P	–
	E	73.1, 21.0, 10.3, 6.0
	M	–
	MM	–
	N	–
CONS.6 – Price trends over next 12 months	*PP*	11.2, 8.4
	P	–
	E	51.2, 8.5
	M	12.0
	MM	–
	N	64.0, 27.7, 12.0, 10.0
CONS.7 – Unemployment expectations over next 12 months	*PP*	48.8, 11.8, 6.0
	P	85.3, 12.0
	E	51.2, 11.8, 4.8
	M	78.8, 12.0
	MM	–
	N	–
CONS.8 – Major purchases at present	*PP*	17.7, 6.0
	E	12.0, 8.7, 4.9, 4.0
	MM	11.6, 6.0, 3.0, 2.0
	N	11.1
CONS.9 – Major purchases over next 12 months	*PP*	–
	P	–
	E	–
	M	23.3
	MM	32.0, 18.6, 13.1, 8.7, 5.5, 3.5
	N	–
CONS.10 – Savings over next 12 months	*PP*	6.0
	P	10.3
	M	4.7, 3.0
	MM	14.0, 8.3, 7.0, 6.0, 3.0
	N	–
CONS.11 – Statement on financial situation of household	*PP*	–
	P	–
	E	–
	M	6.0
	MM	–
	N	–

Notes: Sign – denotes that no cyclical components have been detected or that the detected peaks are of very low energy so as to be considered relevant.

4. Design of a filter for survey indicators

In this section we design a filter for business and consumer survey expectations. The selection criterion of the filtering process is based on the results of the previous section. Our objective is to tailor make a filter to extract the components of the survey indicators with periodicities similar to those that can be found in the dynamics of economic activity.

As noted in Section 2, business and consumer survey expectations are influenced by unpredictable events. This irregular component is considered important in interpreting the data, so balances are constructed using the seasonally-adjusted series instead of the trend. As a result, the EC applies a specific method for seasonal adjustment known as Dainties. While the main advantage of this method is the absence of revisions when adding data at the end of the time series, Dainties is based on a set of assumptions regarding the structure of the components of the series, and applies filters that are decided *a priori* and that do not depend on the data.

To overcome these limitations, in this paper we propose a different approach based on spectral analysis that avoids making any assumption regarding the structure of the components. With this aim, we first estimate the spectral density function of seasonally-adjusted GDP growth in the EA. We use GDP data from the OECD.Stat (Organisation for Economic Co-operation and Development, 2020).

Second, we present the filter, describing its properties. Third, we apply the filter to all business and consumer survey indicators (Table 1). Then, we apply expression (1) to the smoothed industry survey indicators and expression (2) to the smoothed consumer survey indicators in order to compute the balance statistic of the filtered time series. Finally, we assess the performance of the resulting filtered balance statistic (*BF*) by comparing it to the seasonally-adjusted balance (*BS*) published by the EC according to the procedure described in Section 2.

The selection criterion of the filter is based on the spectral analysis of economic activity. We start by computing the spectral density function of seasonally-adjusted monthly GDP growth in the EA. The obtained result is presented in Fig. 6.

In order to extract the components of the survey responses that are closest to those observed in seasonally-adjusted GDP growth (Fig. 6), we have opted for a zero-phase low-pass filter. More specifically, we have chosen a Butterworth filter, which is a type of signal processing filter designed to have a frequency response as flat as possible in the pass band (Butterworth, 1930). As a result, the Butterworth filter is also referred to as a maximally flat magnitude filter (Bianchi and Sorrentino, 2007). See Hamming (1998) and Proakis and Manolakis (1996) for detailed information on the properties of low-pass filtering. In Fig. 6 we overlap the spectral density function of seasonally-adjusted GDP growth to the frequency response of the resulting filter.

With the aim of extracting the components of survey indicators with periodicities closest to those observed in the business cycle, we have applied a Butterworth filter of order 10. As seen in Fig. 6, the estimated spectral density function of GDP shows a sharp decrease in energy around the normalised frequency 0.1, i.e. 10 months. Consequently, the cut-off frequency was fixed at 0.1. The selection of the cut-off frequency is a critical issue, as narrower frequencies would affect the dynamic response of the filter, leaving out components that are significant in terms of amplitude, and which are useful for economic analysis.

As a result of applying the filter, we select the frequency components to the left of the cut-off frequency, which are located in the region that accumulates most of the energy and reject the frequency components of the survey indicators that do not have a counterpart in the GDP time series. The filtering procedure is performed by a finite difference equation of the form:

$$y(t) = \sum_{n=1}^{N} a_n y(t-n) + \sum_{m=0}^{M} b_m x(t-m) \tag{4}$$

where $y(t)$ refers to the output of the filter and, $x(t)$ to the time sequence to be filtered. $\{a_n, b_m\}$ are the coefficients of the filter. The order of the filter is computed as the $Max(N, M)$. See Oppenheim and Schafer (2010) for a detailed analysis of the properties of the filter. The coefficients of the filter are computed by means of the scipy library in Python (https://scipy.org/).

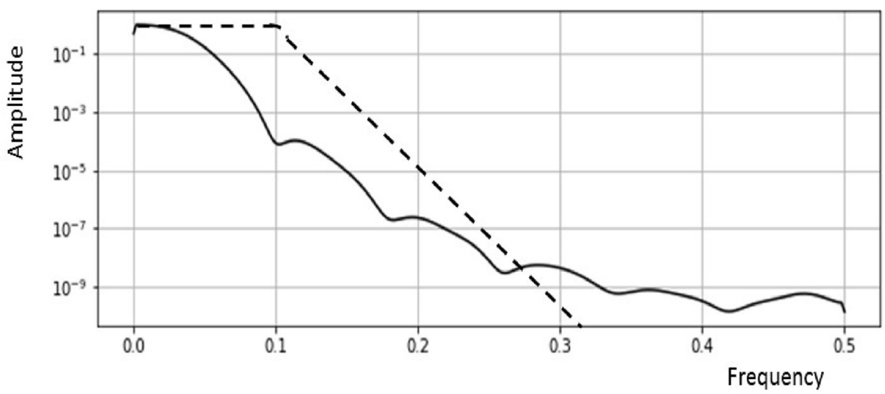

Notes: The continuous line represents the spectral density function of seasonally-adjusted GDP growth. The dashed line represents the frequency response of the low-pass zero-phase filter.

Fig. 6. Spectral density function of seasonally-adjusted GDP growth and frequency response of the low-pass zero-phase filter. Notes: The continuous line represents the spectral density function of seasonally-adjusted GDP growth. The dashed line represents the frequency response of the low-pass zero-phase filter.

Notes: The blue line represents the evolution of the filtered balance (BF), the black dotted line the evolution of the seasonally-adjusted balance (BS) and, the grey line the year-on-year growth rates of seasonally-adjusted monthly GDP for the EA.

(caption on next page)

Fig. 7. Filtered balance (*BF*) vs. seasonally-adjusted balance (*BS*) – Industry survey. Notes: The blue line represents the evolution of the filtered balance (BF), the black dotted line the evolution of the seasonally-adjusted balance (BS) and, the grey line the year-on-year growth rates of seasonally-adjusted monthly GDP for the EA. (For interpretation of the references to colour in this figure legend, the reader is referred to the Web version of this article.)

The zero-phase is obtained by filtering the time series twice: first, from the past values of the series to the future ones, and then, from the future values to the past values. This process allows compensating for the delays introduced by the filtering process. As a result, the resulting filtered survey indicators are temporally aligned with the original time series. Note that this procedure can only be done retrospectively.

Next, we proceed to apply the low-pass filter to all survey indicators (Table 1). Then, we compute the balance statistic of the filtered series. By applying expression (1) to the smoothed industry survey indicators and, expression (2) to the smoothed consumer survey indicators, we obtain a filtered balance statistic (*BF*) for all survey indicators. In order to assess its performance, in Figs. 7–9 we compare them to the seasonally-adjusted balances (*BS*) published by the EC according to the Dainties method described in Section 2. In all graphs we include the year-on-year growth rates of seasonally-adjusted monthly GDP in the EA.

Regardless of how they are processed, survey expectations in the form of balances show a strong and positive relation to economic growth. There is strong evidence in the literature regarding the predictive power of survey expectations. Altug and Çakmakli (2016), obtained superior forecasts of inflation when incorporating survey expectations for Brazil and Turkey. Claveria et al. (2020) have recently found that country-specific sentiment indicators recursively generated by means of genetic programming outperform time series models in predicting quarterly GDP growth rates. Similarly, Juhro and Iyke (2020) showed that accounting for consumer and business sentiments improved forecast accuracy of consumption in Indonesia. Girardi et al. (2016) and Sorić et al. (2019) obtained similar results for the EA.

In Figs. 7–9 it can also be seen that both the filtered balance (*BF*) and the seasonally-adjusted balance (*BS*) are highly correlated, although the evolution of *BF* tends to be smoother than that of the *BS*, especially for the consumer survey indicators. As opposed to the *BS*, which is obtained by applying a seasonal adjustment method dependent on a set of assumptions about the structure of the components of the series, the proposed procedure: i) is specifically designed for this particular set of data, ii) allows extracting the those periodicities similar to the ones that can be found in the dynamics of economic activity, and ii) is directly implementable.

5. Concluding remarks

The main objective of the paper is two-fold. On the one hand, we aim to detect the underlying existing periodicities in business and consumer survey expectations. On the other hand, we aim to provide researchers with a filter especially designed for survey indicators that allows circumventing the application of filters based on *a priori* assumptions regarding the structure of the components.

First, with the objective of giving some insight into the hidden periodicities of business and consumer survey data, we have used the Welch algorithm to compute the specific periodic components of each question and their respective reply options. We have found remarkable differences between business survey indicators and consumer survey indicators. The former, present stronger cyclical components than the indicators form the consumer survey. In the case of the industry survey expectations, we detected a common cyclical component among reply options, corresponding to approximately four to five years, as well as other higher frequency periodicities: twelve, six, four and three months. The cyclical component corresponding to four years was absent in the case of the neutral category of response for the question regarding 'selling price expectations for the months ahead'. For the higher frequency periodicities, the neutral category of response of the industry survey tended to show less marked peaks and of smaller amplitude than the other two extreme response categories.

On the contrary, for consumer survey indicators the intensity of the cycles varied across questions and reply options. In most cases we did not detect any cyclical components, and when we did, these peaks were of very low energy so as to be considered relevant. We also observed that the detected lower frequency periodicities showed a very irregular pattern across questions and reply options.

To a certain extent, these findings could be in part derived from the fact that the wording of the questions differs between both surveys. While business survey indicators mostly refer to the evolution of the variables 'in the recent months' or 'for the months ahead', consumer survey indicators refer to the situation 'over the last twelve months' or 'over the next twelve months'. Additionally, we want to point out the different nature between the questions of both surveys, since the questions of the consumer survey refer to objective variables, while the questions of the business survey refer to specific factors of the firm.

Second, we have designed a filter for business and consumer expectations. The filtering process was aimed at extracting the components of the survey indicators with periodicities similar to those that can be found in the dynamics of economic activity. Therefore, we opted for a low-pass filter that does not need a sharp cut-off frequency and that does not distort the wave form. Specifically, we used a Butterworth filter and applied a zero-phase filtering process to preserve the time alignment of the time series.

To assess the performance of the filtered series, we computed the balance statistic and compared its evolution to that of the seasonally-adjusted balances obtained with the Dainties methodology used by the European Commission. We observed that both statistics showed a highly correlated evolution, although filtered balances tended to be smoother than the seasonally-adjusted ones,

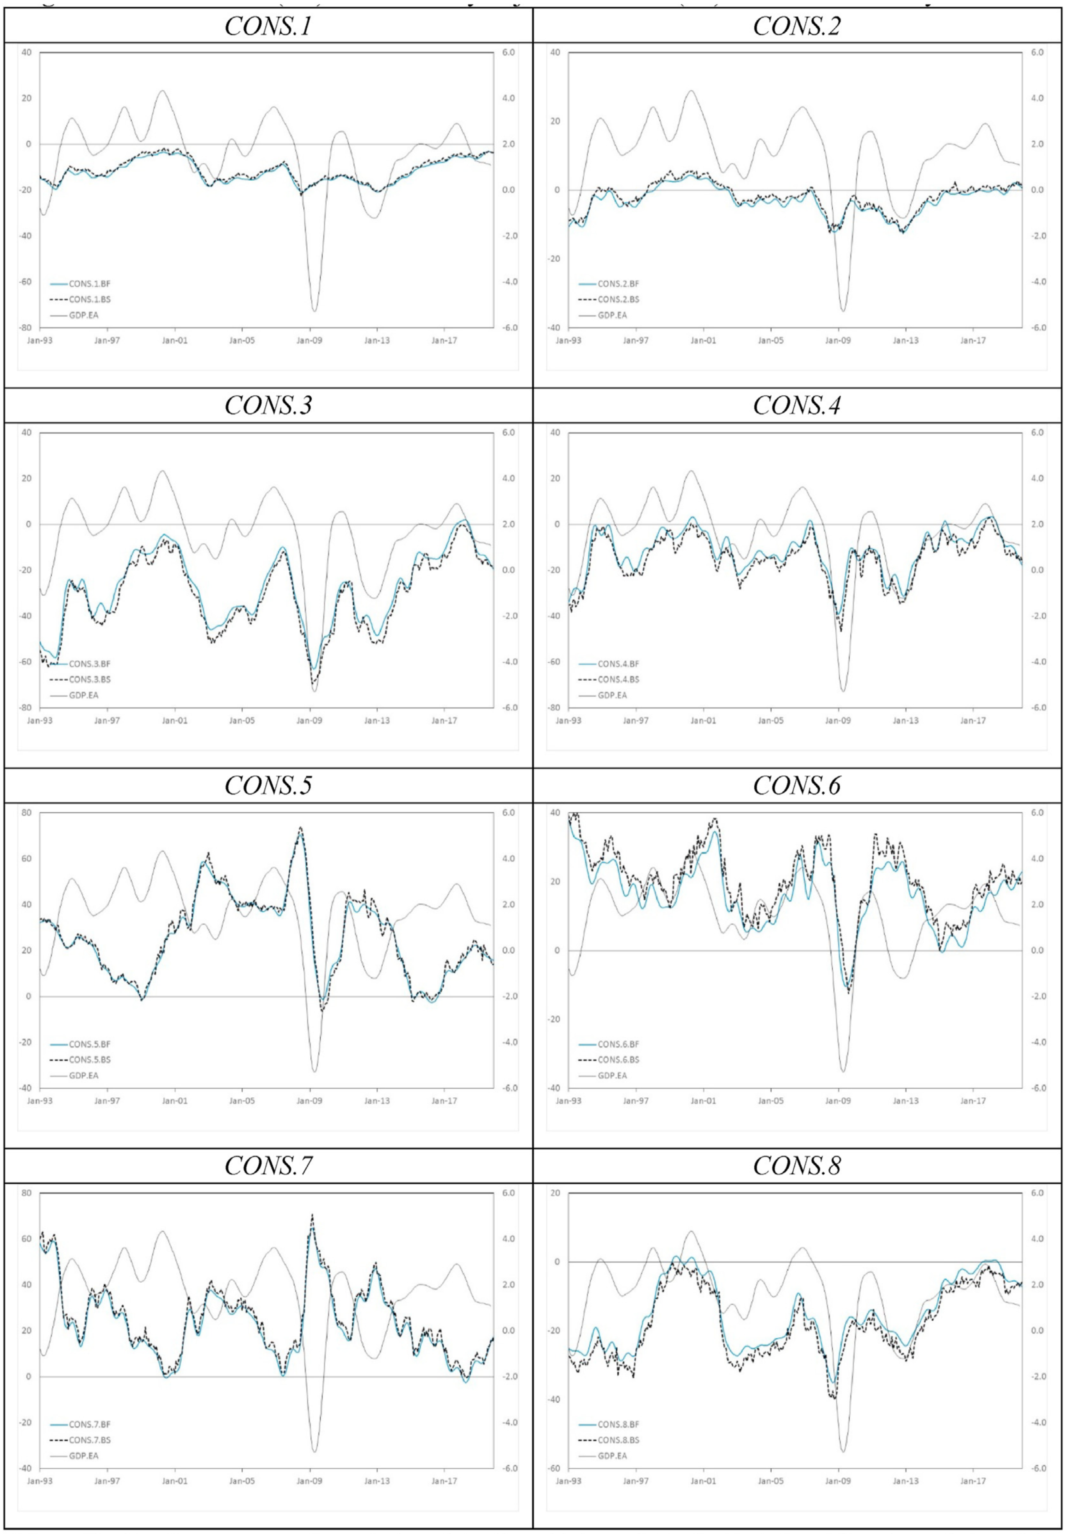

Notes: The blue line represents the evolution of the filtered balance (BF), the black dotted line the evolution of the seasonally-adjusted balance (BS) and, the grey line the year-on-year growth rates of seasonally-adjusted monthly GDP for the EA.

(caption on next page)

Fig. 8. Filtered balance (*BF*) vs. seasonally-adjusted balance (*BS*) – Consumer survey. Notes: The blue line represents the evolution of the filtered balance (BF), the black dotted line the evolution of the seasonally-adjusted balance (BS) and, the grey line the year-on-year growth rates of seasonally-adjusted monthly GDP for the EA. (For interpretation of the references to colour in this figure legend, the reader is referred to the Web version of this article.)

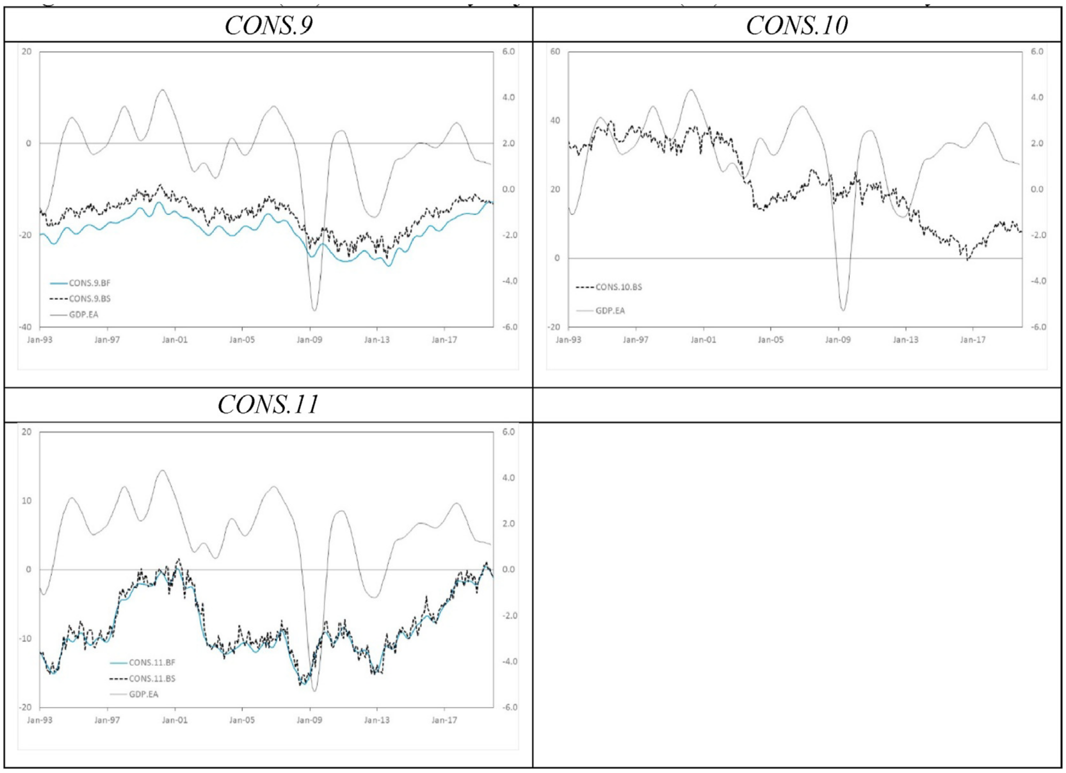

Notes: The blue line represents the evolution of the filtered balance (BF), the black dotted line the evolution of the seasonally-adjusted balance (BS) and, the grey line the year-on-year growth rates of seasonally-adjusted monthly GDP for the EA.

Fig. 9. Filtered balance (*BF*) vs. seasonally-adjusted balance (*BS*) – Consumer survey. Notes: The blue line represents the evolution of the filtered balance (BF), the black dotted line the evolution of the seasonally-adjusted balance (BS) and, the grey line the year-on-year growth rates of seasonally-adjusted monthly GDP for the EA. (For interpretation of the references to colour in this figure legend, the reader is referred to the Web version of this article.)

especially for the consumer survey indicators.

Finally, we want to note that this is a preliminary examination of the periodicities of business and consumer survey indicators by means of spectral analysis. The analysis focused solely on monthly survey data for the Euro Area. One issue left for further research is the extension of the analysis to quarterly data and to the rest of the countries included in the surveys. We also want to point out that one drawback of the proposed approach for the detection of periodicities is its lack of accuracy for components of extremely low frequency. In this sense, a future line of research is the implementation of recently developed signal processing techniques such as Prism. Regarding the filtering of survey indicators, another issue left for further research is the evaluation of alternative filters.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

Altug, S., Çakmakli, C., 2016. Forecasting inflation using survey expectations and target inflation: evidence from Brazil and Turkey. Int. J. Forecast. 32 (1), 138–153.
Bartlett, M.S., 1948. Smoothing periodograms from time-series with continuous spectra. Nature 161, 686–687.
Bartlett, M.S., 1950. Periodogram analysis and continuous spectra. Biometrika 37 (1–2), 1–16.
Bianchi, G., Sorrentino, R., 2007. Electronic Filter Simulation & Design. McGraw Hill Professional, New York, N.Y.
Butterworth, S., 1930. On the theory of filter amplifiers. Wireless Eng. 7, 536–541.
Claveria, O., Monte, E., Torra, S., 2020. Economic forecasting with evolved confidence indicators. Econ. Modell. 93, 576–585. http://scholar.google.es/citations?view_op=view_citation&hl=es&user=j5u1QI8AAAAJ&cstart=20&citation_for_view=j5u1QI8AAAAJ:NXb4pA-qfm4C.
European Commission, 2016. The Joint Harmonized EU Programme of Business and Consumer Surveys. User Guide.
European Commission, 2019. Business and Consumer Surveys. Available at: https://ec.europa.eu/info/business-economy-euro/indicators-statistics/economic-databases/business-and-consumer-surveys_en. (Accessed 6 January 2020).
Gardner, W., 1987. Introduction to Einstein's contribution to time-series analysis. IEEE ASSP Mag. 4 (4), 4–5.
Gelper, S., Croux, C., 2010. On the construction of the European economic sentiment indicator. Oxf. Bull. Econ. Stat. 72 (1), 47–62.
Girardi, A., Gayer, C., Reuter, A., 2016. The role of survey data in nowcasting euro area GDP growth. J. Forecast. 35 (5), 400–418.
Gottman, J.M., 1981. Time-series Analysis. A Comprehensive Introduction for Social Scientists. Cambridge University Press, New York.
Hamming, R.W., 1998. Digital Filters, third ed. Dover Publications, Mineola, N.Y.
Henry, M., 2021. Spectral analysis techniques using Prism signal processing. Measurement 169, 108491.
Juhro, S.M., Iyke, B.N., 2020. Consumer confidence and consumption in Indonesia. Econ. Modell. 89, 367–377.
Kahneman, D., Tversky, A., 1971. Belief in the law of small numbers. Psychol. Bull. 76 (2), 105–110.
Oppenheim, A.V., Schafer, R.W., 2010. Discrete-time Signal Processing, third ed. Prentice Hall, Upper Saddle River, N.J.
Organisation for Economic Co-operation and Development, 2020. OECD.Stat. https://doi.org/10.1787/data-00900-en. Available at: (Accessed 9 January 2020).
Proakis, J.G., Manolakis, D.G., 1996. Digital Signal Processing, third ed. Prentice Hall, Upper Saddle River, N.J.
Sorić, P., Lolić, I., Claveria, O., Monte, E., Torra, S., 2019. Unemployment expectations: a socio-demographic analysis of the effect of news. Lab. Econ. 60, 64–74.
Welch, P.D., 1967. The use of Fast Fourier Transform for the estimation of power spectra: a method based on time averaging over short, modified periodogram. IEEE Trans. Audio Electroacoust. 15 (2), 70–73.