Cycle Identification: An old approach to (relatively) new statistics

Paper presented at the 53rd New Zealand Association of Economists conference,
at Palmerston North, New Zealand,
27 June 2012

Louis Liu, Eldon Paki, James Stonehouse, and Jing You

Development, Environment, Satellites, and Strategy,
Statistics New Zealand
PO Box 2922
Wellington, New Zealand
info@stats.govt.nz
www.stats.govt.nz

Liability
The opinions, findings, recommendations, and conclusions expressed in this paper are those of the authors. They do not represent those of Statistics New Zealand, which takes no responsibility for any omissions or errors in the information in this paper.

Citation
Liu, L, Paki, E, Stonehouse, J, and You, J (2012, June). *Cycle identification: An old approach to (relatively) new statistics.* Paper presented at the 53rd New Zealand Association of Economists Conference, Palmerston North.

Abstract

Whether looking at New Zealand's economy or the global economy, it is important to understand and appreciate business and growth cycles. The first step to understanding cycles is to correctly identify them. Economists usually reference the time domain to identify cycles. We take an alternative view, and reference the frequency domain instead.

We ran Fourier analysis on several data sources to demonstrate how the frequency domain can reveal cyclical behaviour. The data included electricity demand, foreign exchange, monthly retail sales, quarterly GDP, labour market, and productivity statistics.

Introduction

The importance of identifying business cycles

There are two reasons to consider and understand business and growth cycles when viewing any economy. The first is that business cycles are an intrinsic part of a nation's macroeconomy. Therefore, understanding business cycles helps us to understand that nation's macroeconomic performance. The second is the impact at the micro level – business cycles affect individuals and firms.

Macroeconomists contend the study of business cycles is valuable to society. In 2004, the Nobel Prize in Economics went to Finn E. Kydland and Edward C. Prescott for their contribution toward understanding the drivers of business cycles.

Economists usually use the time domain to identify cycles

Economists have done a lot of work on identifying business cycles in the time domain. This is due in part to the data. Macroeconomic data is time series data, which means the data is naturally thought of and visualised in the time domain.

It has therefore been natural for the techniques of identifying cycles to develop within the time domain framework. For example, the algorithm approach of Bry and Boschan (1971) in identifying cycle turning points was illustrated using macroeconomic data in the time domain. Their first illustration used the US unemployment rate for the years 1929 to 1965.

Other disciplines use the frequency domain to identify cycles

Rather than use the time domain to conduct their analysis of the data, analysts in other fields use the frequency domain. For example, electrical engineers routinely analyse time series data from their field to identify its cyclical components. They focus on identifying the cycle, its frequency, and its amplitude from the data.

We use the frequency domain on economic-related data

Our paper follows the engineering approach by using the frequency domain to analyse time series data. We use data from Statistics New Zealand and data from the Electricity Authority. We also use gross domestic product (GDP) data from Hall and McDermott (2011) and foreign currency data from Interbank FX.

Our objectives

We have two objectives in this paper. Our first objective is to show how to use the frequency domain on data with known cyclical behaviour. We do this by using quarterly sales data and electricity data. Our second objective is to use it on data with less certain cyclical or periodic behaviour. In this category lies the other time series data that we investigate.

The next section reviews the work that Statistics NZ has already done on cycles. The section following that provides a brief introduction to the theory behind the frequency domain. The section after that describes the data and data sources used in this paper. After this, we present our results. We then summarise our results at the end of the paper.

Statistics NZ's previous research on cycles

In a 2007 paper, Extracting growth cycles from productivity indexes, Statistics NZ described their research on identifying cycles in productivity statistics. The department wanted to minimise the impact of factors that vary within a cycle so that people using the data could compare productivity performance between periods.

Three methods for identifying the trend in productivity data

In Extracting growth cycles from productivity indexes Statistics NZ presented three methods for identifying the trend in productivity data: the Hodrick-Prescott filter, the Baxter-King filter, and the Henderson filter. Each of the three filters was used to remove the trend, or more precisely, leave behind the deviations from trend. Statistics NZ defined a peak as the highest deviation above the trend and a cycle as the time elapsed between one peak the next in time series data.

How to use the Hodrick-Prescott filter to remove trends in data

The first method used the Hodrick-Prescott filter. When they published their paper on the this filter, Hodrick and Prescott (1997) considered that a time series Y_t can be decomposed into two components; the trend component τ_t, and the cyclical component c_t. That is,

$$Y_t = \tau_t + c_t \quad t = 1,2, \dots, T.$$

Given an appropriately chosen positive value of λ they find the trend component τ_t that solves the following equation

$$\min\left\{\sum_{t=1}^{T}(y_t - \tau_t)^2 + \lambda\sum_{t=2}^{T}[(\tau_{t+1} - \tau_t) - (\tau_t - \tau_{t-1})]^2\right\}.$$

The parameter λ is the smoothing parameter. That is, you can smooth the trend component τ_t by choosing a sufficiently large λ value. In their paper the authors recommend a λ value of 1600 for quarterly data.

Subtracting the Hodrick–Prescott trend component τ_t from time series Y_t leaves the cyclical component c_t. That is,

$$c_t = Y_t - \tau_t \quad t = 1,2, \dots, T.$$

Statistics NZ labelled the cyclical component c_t as the deviation from trend.

How to use the Baxter-King filter to remove trends in data

The second method Statistics NZ investigated uses the Baxter-King filter. In contrast to the Hodrick-Prescott filter, Baxter and King (1999) considered Y_t can be decomposed into three components: the trend component τ_t, the business-cycle component c_t, and the irregular component I_t. That is,

$$Y_t = \tau_t + c_t + I_t \quad t = 1,2, \dots, T.$$

The Baxter-King filter removes the trend component τ_t and the irregular component I_t, leaving the business-cycle componentc_t. That is,

$$c_t = Y_t - \tau_t - I_t \quad t = 1,2, \dots, T.$$

Statistics NZ labelled the business-cycle component c_t as the deviation from trend.

How to use the Henderson filter to remove trends in data

The third method uses the Henderson filter. Henderson derived this filter in 1916 for actuarial purposes. He considered a time series Y_t to be decomposed into a cubic polynomial trend component $P_3(t)$, and a random component c_t, with the property that the expected value of the random component c_t equals zero. That is,

$$Y_t = P_3(t) + c_t \quad t = 1,2, \dots ,T \text{ with } E(c_t) = 0.$$

Henderson estimated the trend component $P_3(t)$ by using weighted least squares methods.

Subtracting the Henderson trend component $P_3(t)$ from the time series Y_t leaves the random component c_t. That is,

$$c_t = Y_t - P_3(t) \quad t = 1,2, \dots ,T.$$

Statistics NZ labelled the random component c_t as the deviation from trend.

There is another way to identify cycles

Rather than use the methods explained above to remove trends and identify a cycle, or the algorithm approach of Bry and Boschan (1971), we used the frequency domain approach. This approach is known as Fourier analysis, and is commonly used in fields such as engineering, geophysics, and metrology. We now provide a brief introduction to Fourier analysis.

Using the frequency domain to identify cycles

For this section, assume the time series data is detrended data. Having graphed the time series Y_t, suppose there is evidence of periodic behaviour. Then a natural model for the time series is

$$Y_t = Rcos(\omega t + \theta)$$

where R is the amplitude, ω is the frequency, and θ is the phase. This is a very simple model in the frequency domain with only one frequency.

Explanation of the harmonic model

In practice, a time series may have several different frequencies. For example, retail sales data may have weekly, monthly, quarterly, and yearly frequencies. In other words, the data may show high, medium, and low frequencies. A natural extension of the model above that allows several frequencies is

$$Y_t = \sum_{j=1}^{k} R_j cos\left(\omega_j t + \theta_j\right).$$

Making use of the identity

$$cos(\omega t + \theta) = cos\omega t cos\theta - sin\omega t sin\theta$$

we obtain

$$Y_t = \sum_{j=1}^{k} \left(a_j cos\omega_j t + b_j sin\omega_j t\right)$$

where $a_j = R_j cos\theta_j$ and $b_j = -R_j sin\theta_j$. This is known as the **harmonic model**.

Explanation of Fourier analysis

To simplify matters we work with frequencies in a range from $-\pi$ to π. Without loss of generality, suppose that $\omega_k = \pi$. Using the identity

$$cos\omega + isin\omega = e^{i\omega}$$

we rewrite the model as

$$Y_t = \sum_{j=-k}^{k} z_j e^{i\omega_j t}$$

where $z_{-j} = \frac{1}{2}(a_j + ib_j)$ and $z_j = \frac{1}{2}(a_j - ib_j)$ are complex numbers for $-k < j < k$.

Using the form in Janacek (2001), the function

$$I(\omega) = \frac{1}{N}\left|\sum_{t=1}^{N} Y_t e^{-i\omega t}\right|^2$$

is known as the periodogram of the time series with $I(\omega)$ called the power of the periodogram at value ω. Its close relative

$$J(\omega) = \frac{1}{\sqrt{N}}\sum_{t=1}^{N} Y_t e^{-i\omega t}$$

is known as the discrete Fourier transform of the time series. The discrete Fourier transform is easily implemented and is part of the Data Analysis package in Microsoft Excel.

The key point of interest with the periodogram is that when you graph $I(\omega)$ with ω along the x-axis, any large peak of the periodogram identifies cyclical behaviour.

Studying a periodogram to identify periodic behaviour is referred to as conducting **Fourier analysis**.

For the interested reader an excellent introduction to the topic is Chatfield (2004).

Example of how to conduct Fourier analysis

We now illustrate how to use Fourier analysis with a stylised example in three parts.

The first part of the example is

$$Y_t = cos(8\pi t)$$

where the amplitude R is one, the frequency ω is 8π, and the phase θ is zero. Figure 1 displays the time series and its corresponding periodogram. The peaks on the time series occur at a frequency of 0.25.

Figure 1

Example part 1

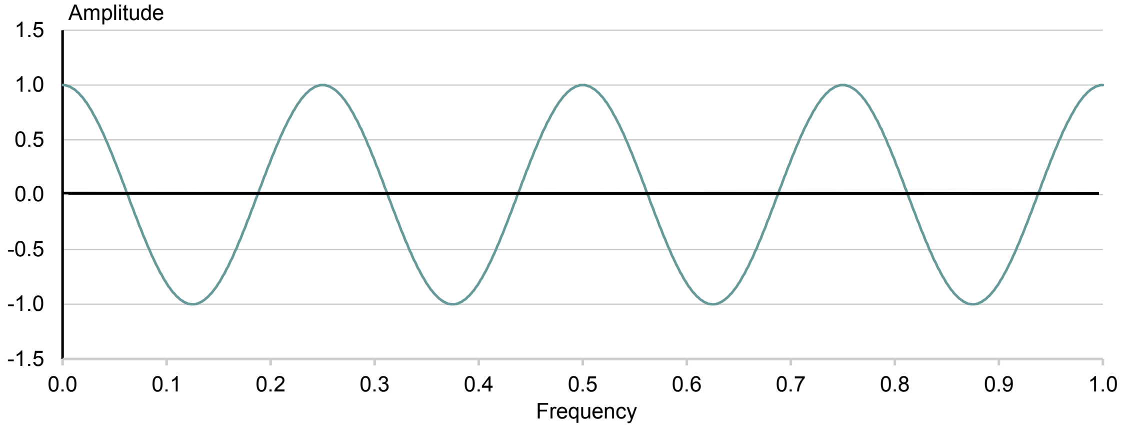

Source: Authors' calculations

Periodogram of example part 1

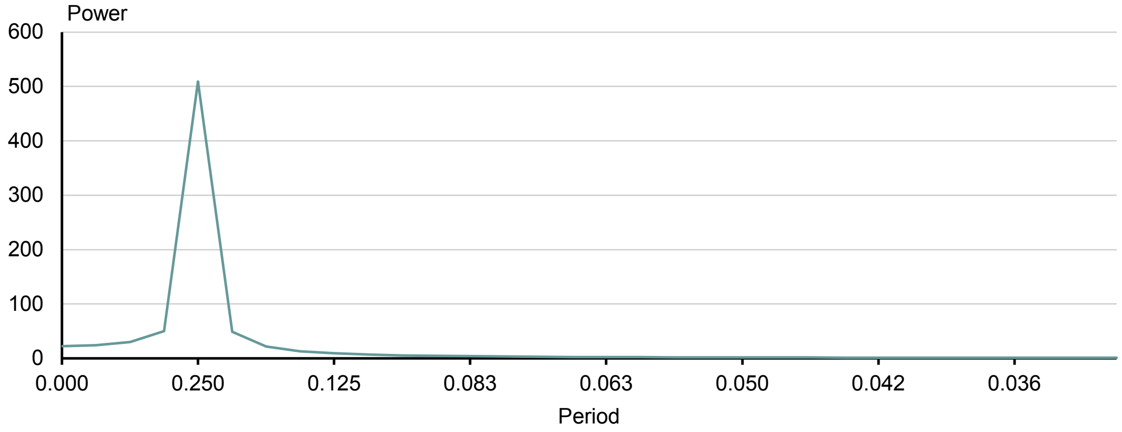

Source: Authors' calculations

The periodogram has a significant peak at 0.250, which confirms the cycle you can see in the time series.

The second part of the example is

$$Y_t = \frac{1}{3}cos(24\pi t)$$

with an amplitude of $\frac{1}{3}$, frequency of 24π, and phase θ of zero. Figure 2 displays the time series and corresponding periodogram.

Figure 2

Example part 2

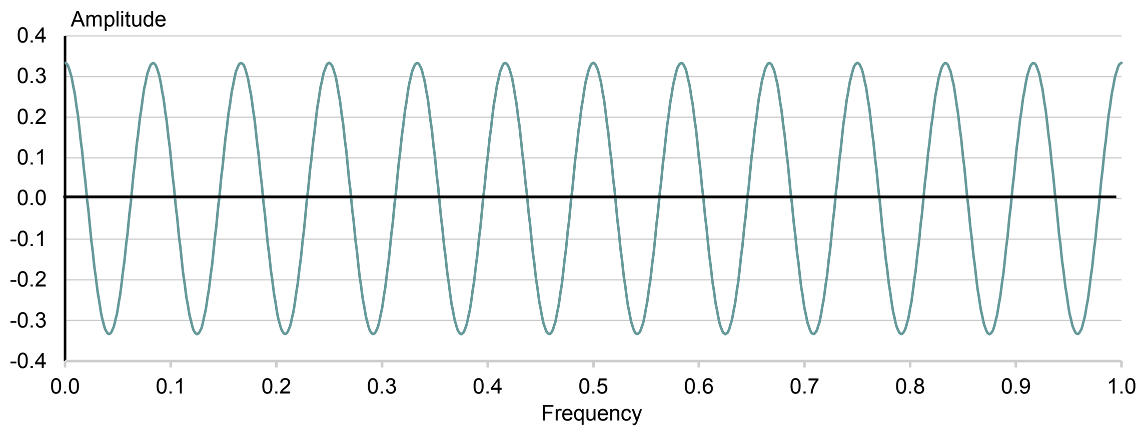

Source: Authors' calculations

Periodogram of example part 2

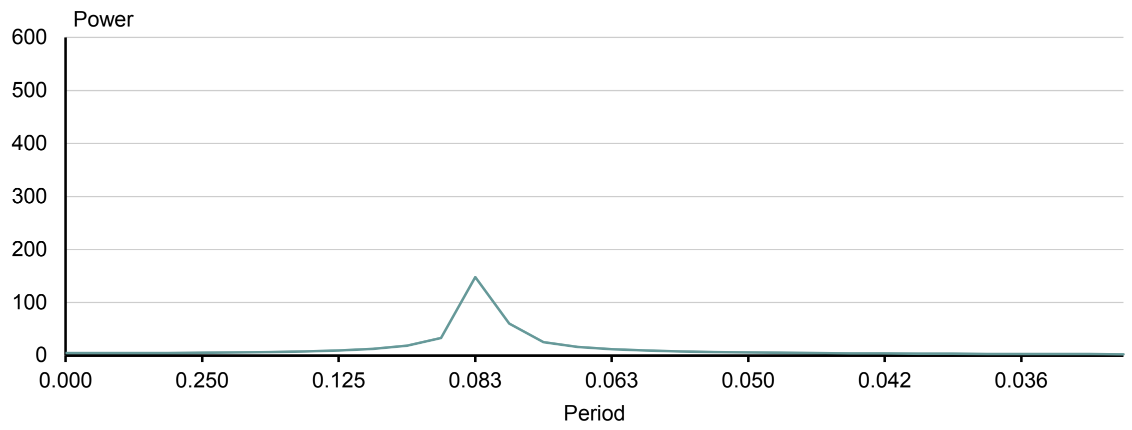

Source: Authors' calculations

The periodogram shows a significant period at 0.083 as expected. Note the shorter upward spike compared with the first periodogram – this is because of the relatively smaller amplitude.

The third part of the example combines the two series:

$$Y_t = cos(8\pi t) + \frac{1}{3}cos(24\pi t).$$

Figure 3 displays the combined time series and associated periodogram.

Figure 3

Example part 3

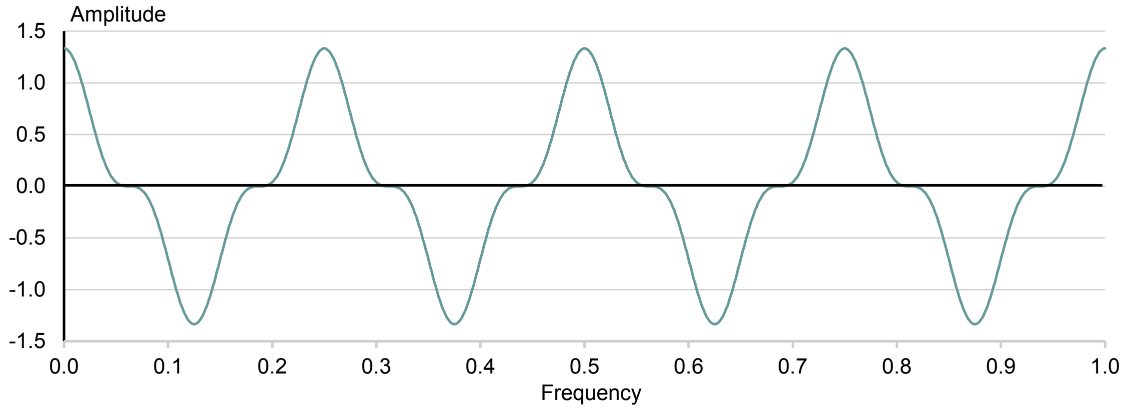

Source: Authors' calculations

Periodogram of example part 3

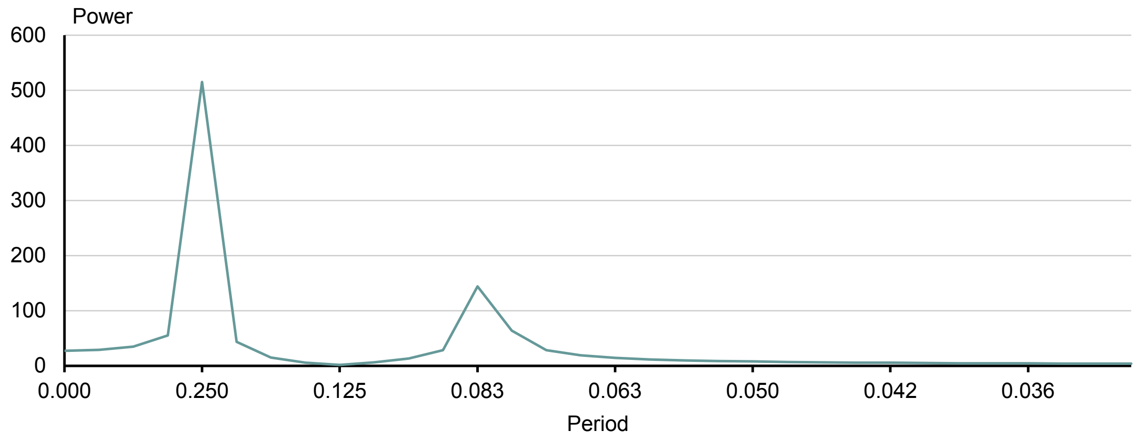

Source: Authors' calculations

The resulting periodogram of the combined series shows the two significant periods at 0.250 and 0.083. When the two frequency components of the new wave form appear together in one periodogram, the frequency and amplitude of each is preserved.

How we used Fourier analysis to identify cycles

We used Fourier analysis to identify the length of a cycle. Bry and Boschan (1971) outlined an approach that identifies cycle start- and end-points. The length of the cycle is inferred from the start- and end-points. Our approach is the reverse of Bry and Boschan's – we identify the length of a cycle, and from that can infer the start- and end-points of a cycle.

Data sources analysed in our paper

We collected data from several sources. The first is electricity data from the Electricity Authority. It is half-hourly data from the Transpower grid exit point at Upper Hutt from 1 January 2010 to 31 December 2011. There are 48 data points for a day and 17,520 data points for a year – a total of

35,040 data points over the two-year period. The units are kilowatt-hours. The second is foreign currency data from Interbank FX. The data is the daily spot price highs and spot price lows of the NZD/USD pair in the period April 2004 to December 2010 – some 2,000 data points. The third is quarterly GDP data from Hall and McDermott (2011). The rest of the data used in our paper is publicly available from Statistics NZ.

Cycles in data with well-defined periodic characteristics

We began by looking at data with well-defined periodic characteristics: retail sales and electricity use.

Retail sales data

We used a periodogram to identify the growth rate cycles in retail sales data. We used monthly retail sales for the Auckland region from Statistics NZ. The data is millions of dollars in current prices, excluding GST. The time period the data covers is May 1995 to September 2010.

Figure 4 shows the retail sales data, retail growth rates, and periodogram of the growth rates. The main feature from the sales data is the spike in retail sales that appears in December of every year. The main feature of the monthly growth rates is the high growth rate in December of every year that is followed by a negative growth rate in January of every year.

Figure 4

Auckland retail sales
May 1995–September 2010

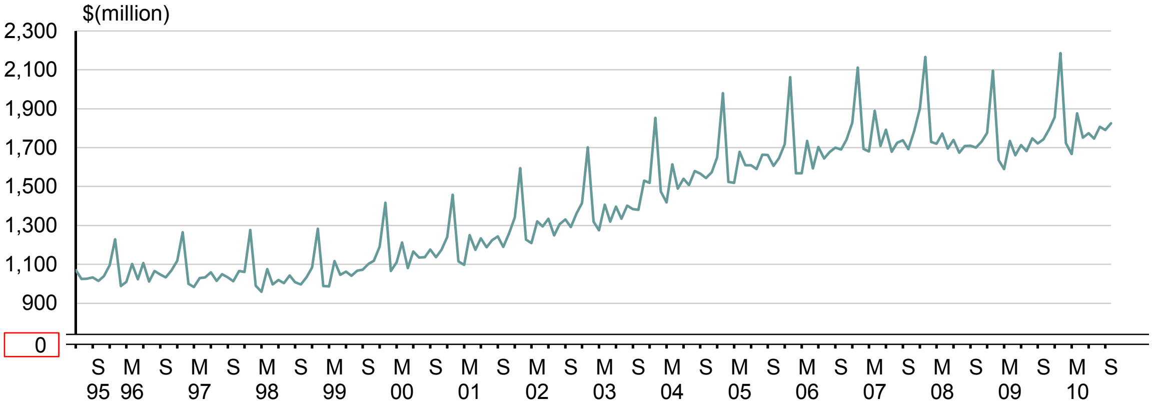

Source: Statistics New Zealand

Auckland retail sales growth rates
June 1995–September 2010

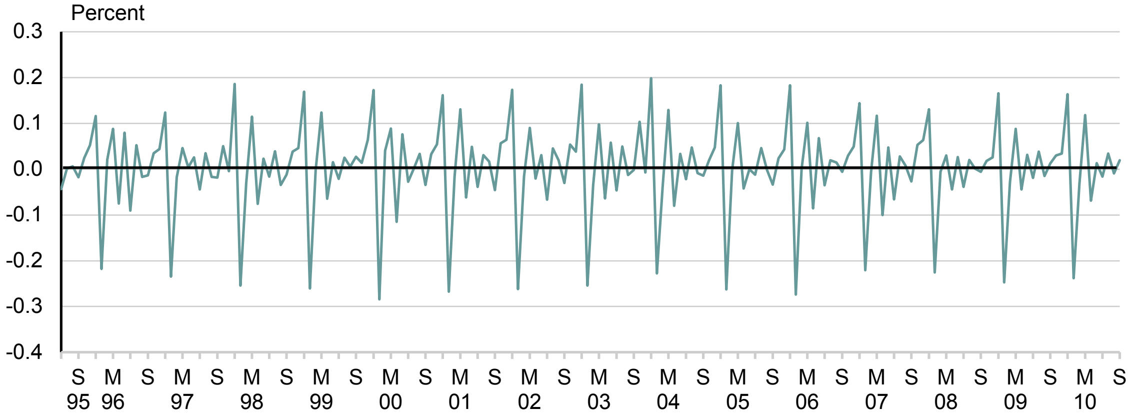

Source: Authors' calculations

Periodogram of Auckland retail growth rates
June 1995–September 2010

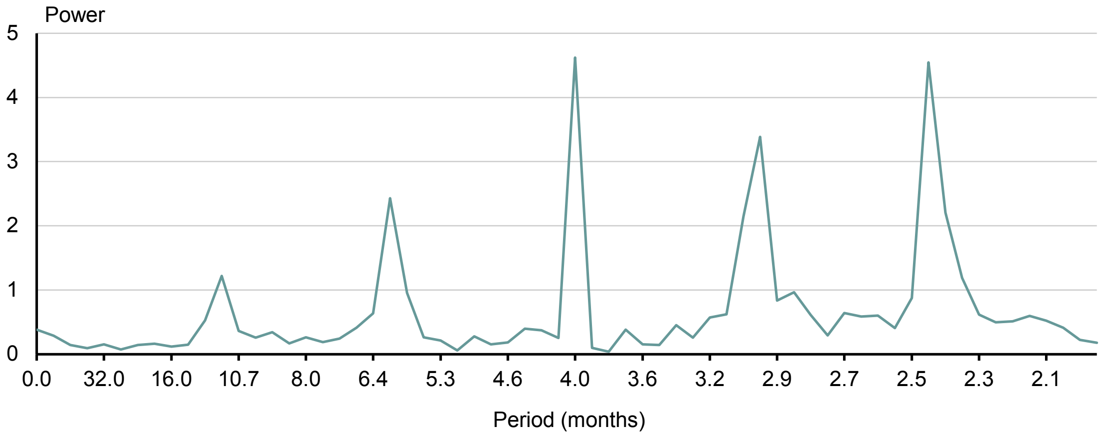

Source: Authors' calculations

The periodogram of the growth rates has five distinct peaks rather than one, which requires an explanation. The leftmost peak shows a period of approximately twelve months, indicating the yearly cycle as expected. The second peak is at six months, or twice the frequency of the first peak. The third peak is at three times the frequency of the first peak. The fourth peak is four times, and the fifth peak is five times the frequency of the first. This periodogram is therefore a harmonic series.

A harmonic series has a unique feature: other frequencies are multiples of the fundamental, or initial, frequency. Regular or periodic patterns show harmonics as seen in this example.

Electricity data

In the next study we show the electricity data from the Electricity Authority. For this example we have combined the half-hourly readings into 730 daily readings and show the daily readings over the two-year calendar period 2010–11. Figure 5 shows the data and periodogram.

The main feature of the data is the clear yearly cycle on display. Electricity use starts relatively low at the beginning of 2010. As the season changes from summer to autumn to winter, electricity use increases. Electricity use peaks mid-way through 2010 then declines as the season changes to spring then to summer. The pattern repeats in 2011. Note the two sharp peaks in 2011 – these indicate Upper Hutt endured two separate cold snaps in 2011.

Less clear is the other cycle in the data. Local minima occur on a weekly basis. Compared with electricity use during weekdays, electricity use on the weekends is lower. The low for the week is always on Saturday or Sunday.

Figure 5

Upper Hutt daily energy use
2010–2011

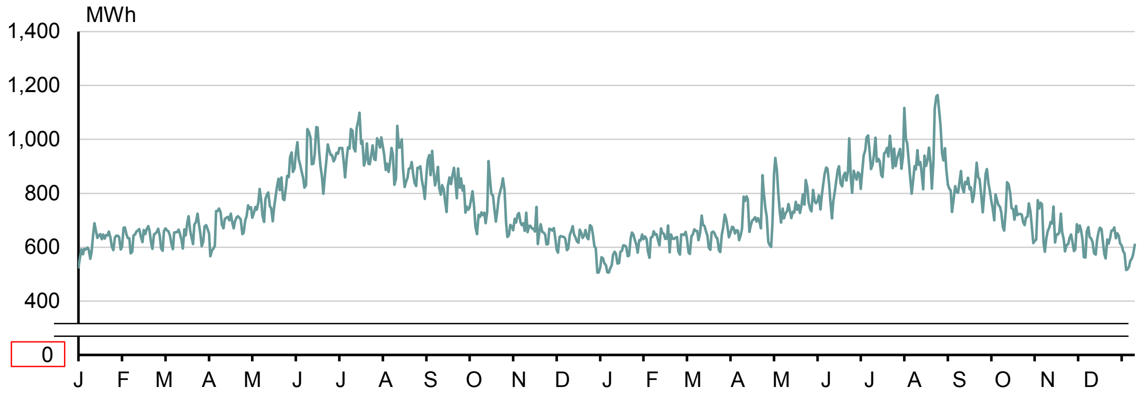

Source: Electricity Authority

Periodogram of Upper Hutt daily energy use
2010–2011

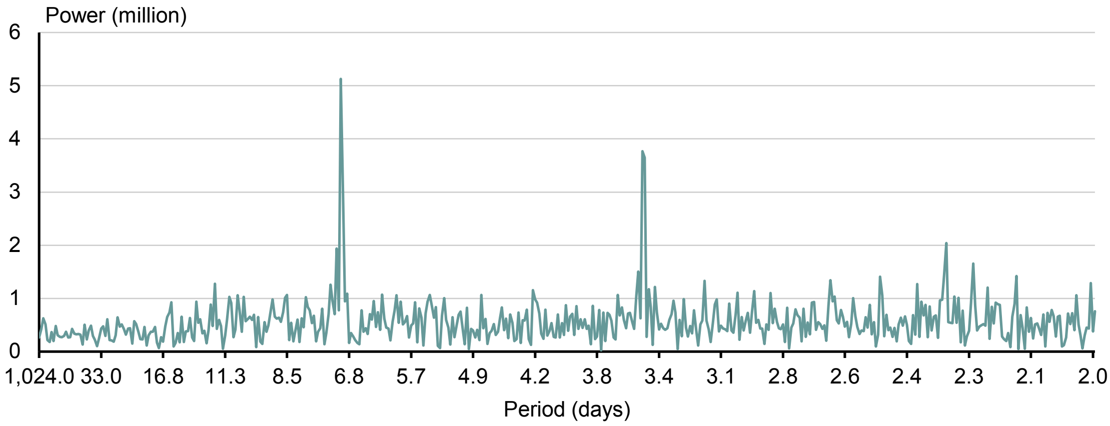

Source: Authors' calculations

Once the data is detrended, it loses its yearly cyclical characteristic and leaves the weekly characteristic. With this in mind, we expect to see a peak in the periodogram of the detrended data that corresponds to the weekly cycle. As the periodogram in Figure 5 illustrates, the period of seven days is evident, as is the second harmonic with a period of three-and-a-half days.

Cycles in data with less certain periodic characteristics

We also looked at data with periodic characteristics that are less certain. That is, we expected the data to have periodic characteristics but the empirical evidence may not support our expectations. The labour market, currency trading, and GDP are the data in this category.

Labour market data

We now consider the labour market. The data is from Statistics NZ's Quarterly Employment Survey and covers the period from the first quarter in 1989 to the first quarter in 2012. The measured sector is a combination of the industries that operate in a competitive or market environment.

Figure 6 displays the data and its periodogram. Employee counts trend downward at the beginning of the series until 1992. Then for the middle part of the 1990s employment counts increase. From mid-1998 to the turn of the century employment is static. Then it resumes its upward trend and peaks in the fourth quarter of 2007.

Figure 6

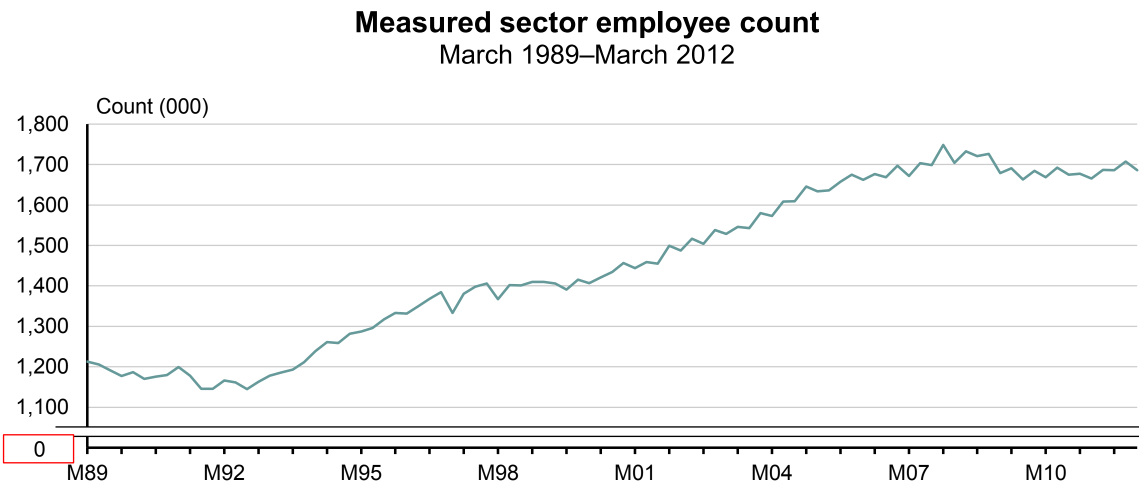

Source: Statistics New Zealand

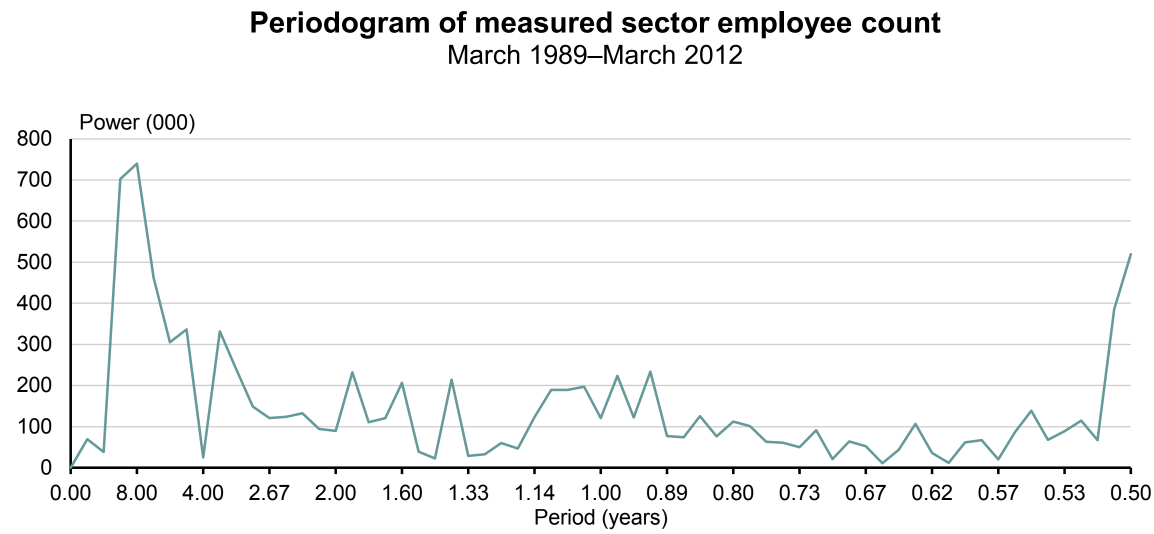

Source: Authors' calculations

The periodogram suggests that for employment in competitive industries, there is a cycle of eight years. This matches the time period for a typical business cycle.

Currency trading data

We investigated daily volatility of the NZD/USD pair spot price from Interbank FX. Volatility in this instance is defined as the difference between the spot price high for the day and the spot price low for the day. The data covers the period from April 2004 to December 2010.

Figure 7 shows the data and its periodogram. Note the increase in volatility mid-August 2007 and the larger volatility increase in October 2008. Prior to 2007 the average daily volatility was $0.0071, or 71 pips. As a contrast, for two consecutive days in August 2007 the volatility was 331 and 345 pips respectively. For the period October 2008, there were three days of extreme volatility. The first movement was 400 pips, the second was 474 pips, and the third was 444 pips.

Figure 7

NZD/USD daily spot price volatility
April 2004–December 2010

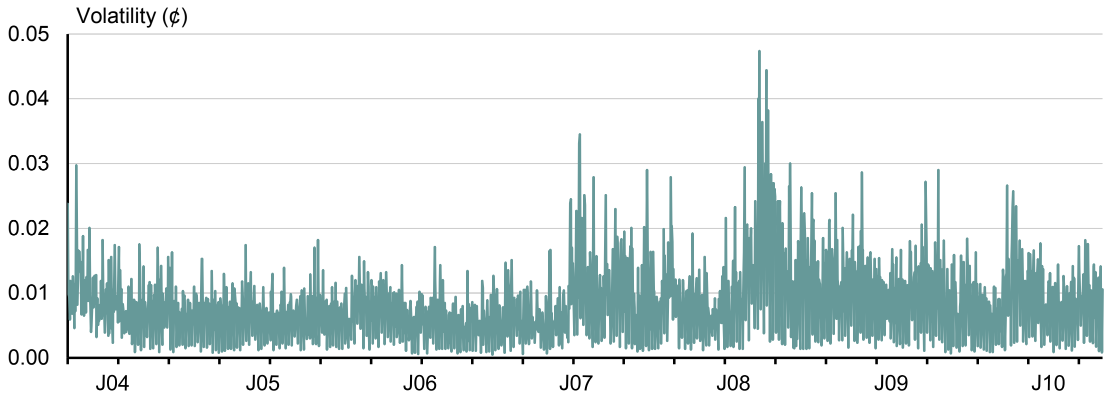

Source: Interbank FX

Periodogram of NZD/USD daily spot price volatility
April 2004–December 2010

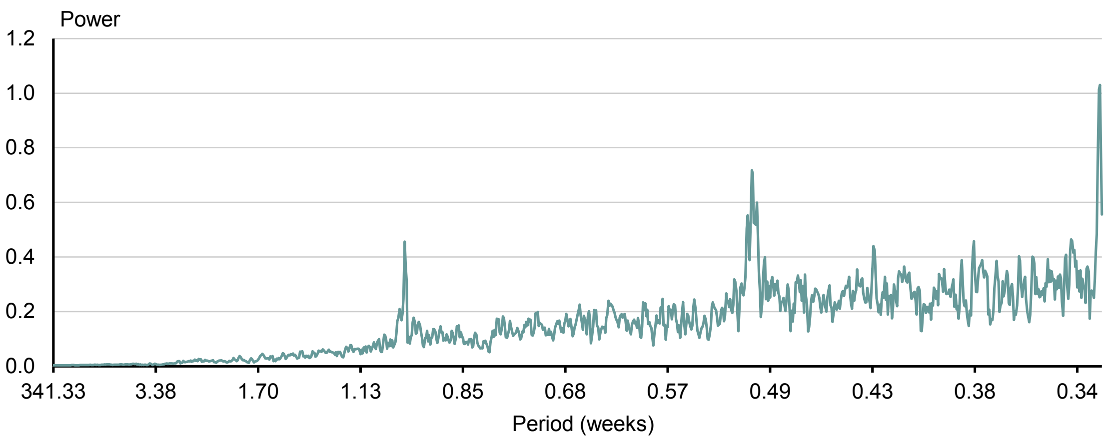

Source: Authors' calculations

The periodogram identified a weekly cycle and its second harmonic. After a careful examination of the data, we found that volatility on the last trading day of the week, in general, is higher than on other trading days of the week. We suggest a reason for this may be that forex traders are less willing to hold open positions over the weekend compared with overnight. So a higher level of trading takes place on the last trading day of the week, with a corresponding increase in the volatility of the spot price.

GDP data

We now consider the GDP data from Hall and McDermott (2011). Figure 8 displays the data and its periodogram. To construct the periodogram we first transformed the GDP data using natural logarithms, detrended the transformed data using the Hodrick-Prescott filter, and conducted a Fourier analysis of the detrended transformed data.

Figure 8

Quarterly real GDP(1)
June 1947–September 2008

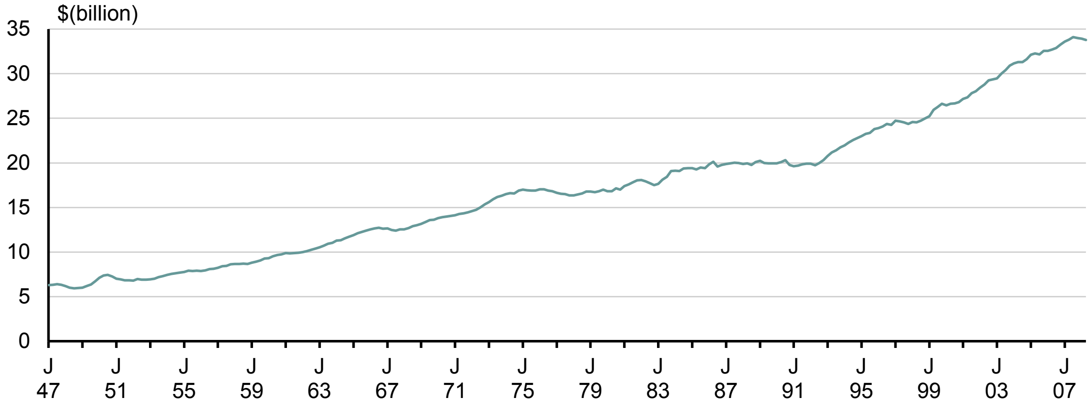

Source: Hall and McDermott (2011)
1. Seasonally adjusted, 1995–96 prices

Periodogram of quarterly log (GDP)
1947–2008

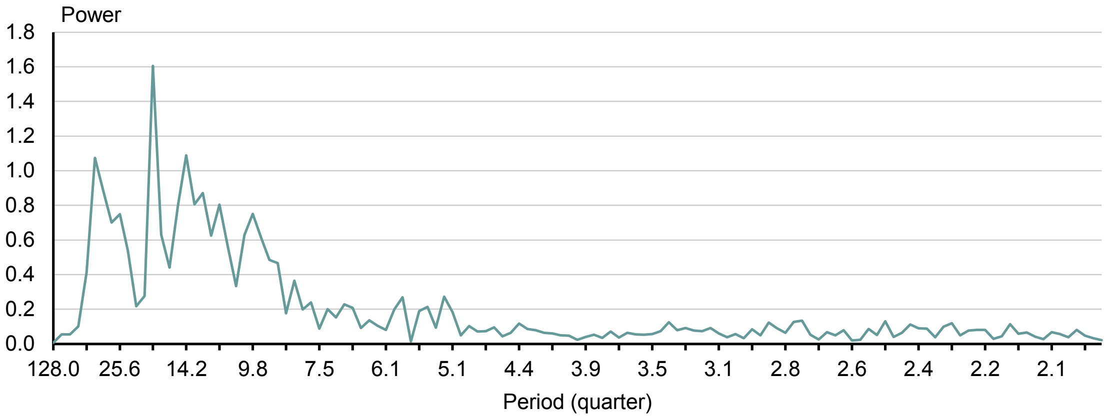

Source: Authors' calculations

The leftmost peak of the periodogram has a power value of 1.07 at 32 quarters, corresponding to a cycle of eight years. The next peak has a power value of 1.60 at approximately 18 quarters, corresponding to a cycle of four-and-a-half years. Because of the spikes in the periodogram we suggest a cycle is present in the data. In addition, because of the distance between the two spikes of 32 quarters and 18 quarters we further suggest that the cycle length varies, and that it varies between 32 and 18 quarters, or four-and-a-half to eight years.

Cycles in data with uncertain periodic behaviour

We now turn our attention to data with uncertain periodic behaviour.

Labour productivity data

We investigated labour productivity statistics for the manufacturing industry. In this and subsequent studies, to construct the periodogram we detrended the data using the Hodrick-Prescott filter and conducted a Fourier analysis on the detrended data.

Figure 9 shows the labour productivity index and associated periodogram. Note the overall upward trend in the index with its peak in 2006.

Figure 9

Labour productivity index: manufacturing
1978–2010

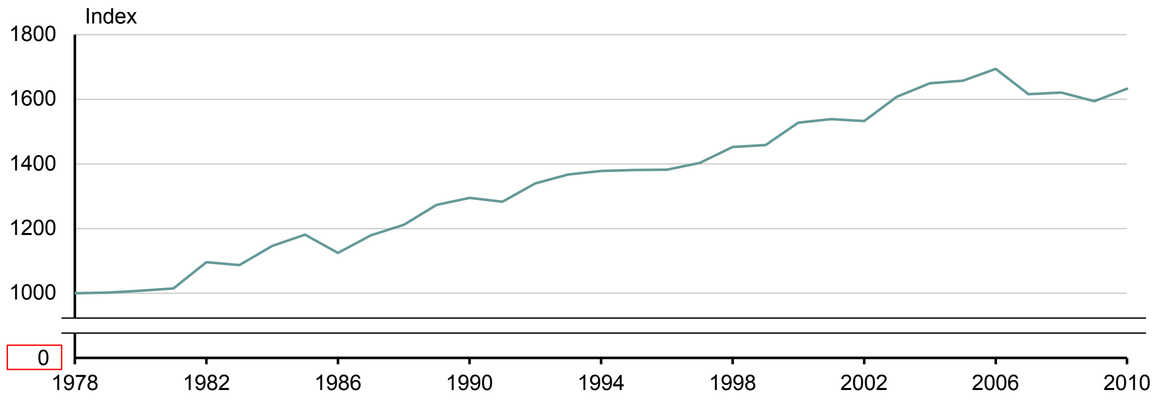

Source: Statistics New Zealand

Periodogram of labour productivity index: manufacturing
1978–2010

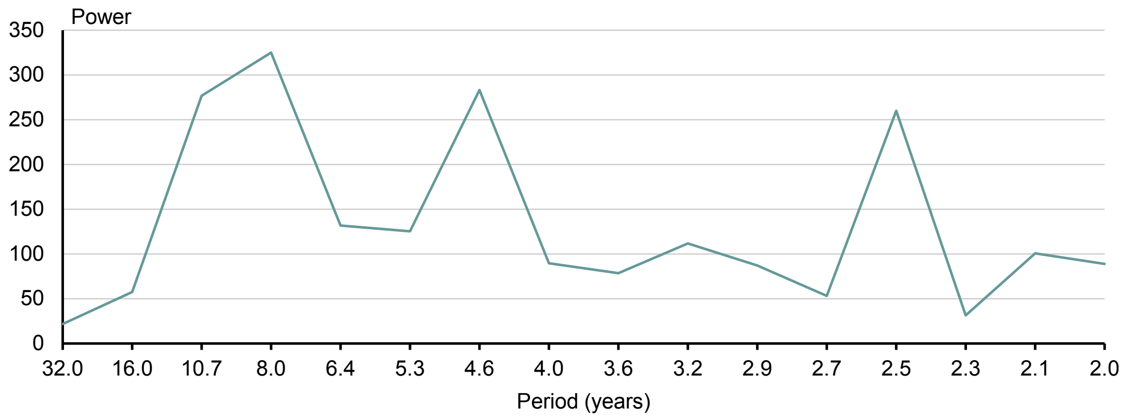

Source: Authors' calculations

In this study the periodogram shows a labour productivity cycle of eight years in the manufacturing industry. That is, labour productivity in the manufacturing industry goes through an eight-year cycle. This result is in accord with our previous analysis of the Quarterly Employment Survey data, which showed employee counts in the measured sector also experienced an eight-year cycle.

Capital productivity input index

We chose to analyse capital productivity input statistics for the mining industry in part because mining is a capital-intensive industry.

Figure 10 displays the data with its periodogram. The data shows a gradual increase from the late 1970s to its smooth peak in 1993. Capital productivity input then gently declines until it bottoms in 2002. Since then, capital productivity input has increased and is starting to show signs of slowing down.

Figure 10

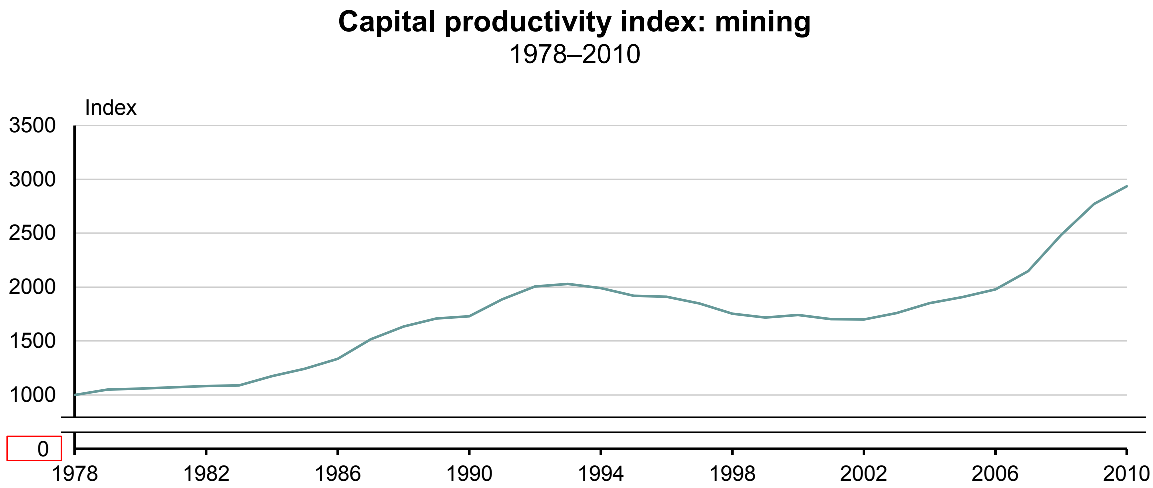

Source: Statistics New Zealand

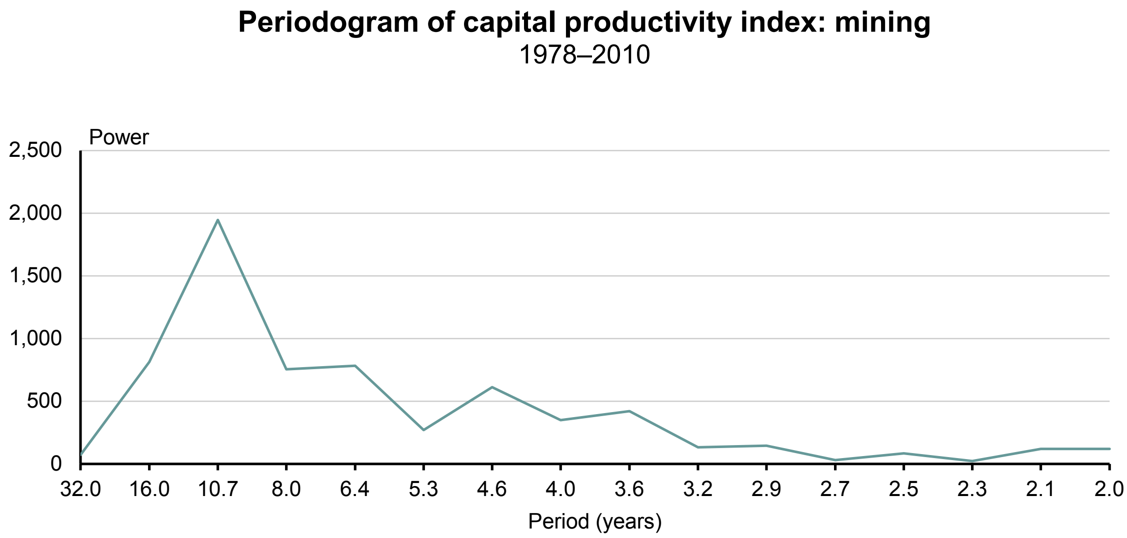

Source: Authors' calculations

The periodogram shows that capital productivity input has a cycle of approximately 11 years. We suggest this length of time for a cycle is reasonable. Capital-intensive industries like mining require large investments, so it would make sense that the cycle would be longer than the typical five to eight years.

Multifactor productivity index data

The last study we conducted used multifactor productivity (MFP) statistics for the New Zealand electricity, gas and water supply industry.

Figure 11 displays the MFP index data and its periodogram. In contrast to the previous two studies, the data here is less smooth and is a story with two parts. The first is the general upward trend until it peaks in 1996. The second is the general downward trend since the peak in 1996. We note that the electricity industry went through major reform in the latter half of the 1990s.

Figure 11

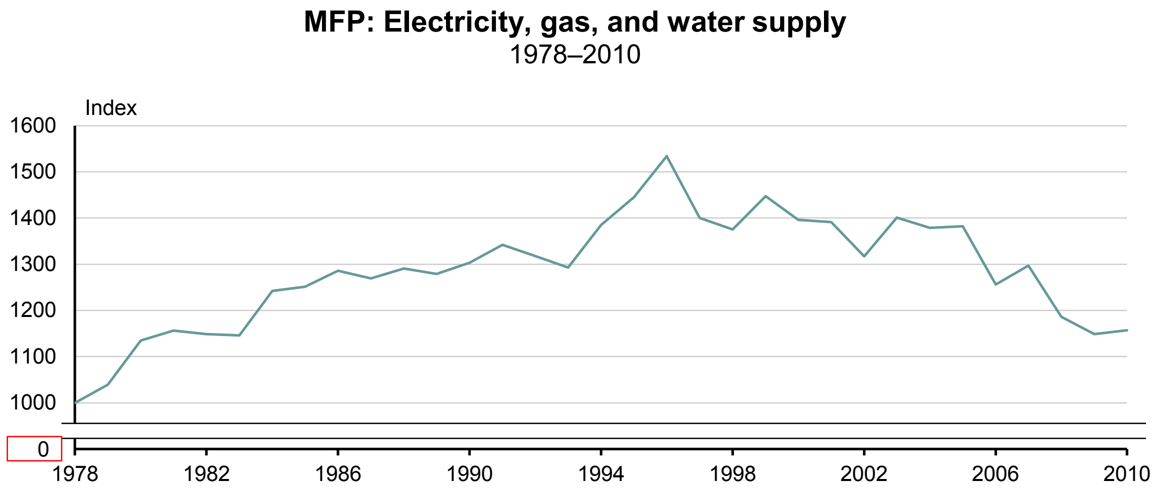

Source: Statistics New Zealand

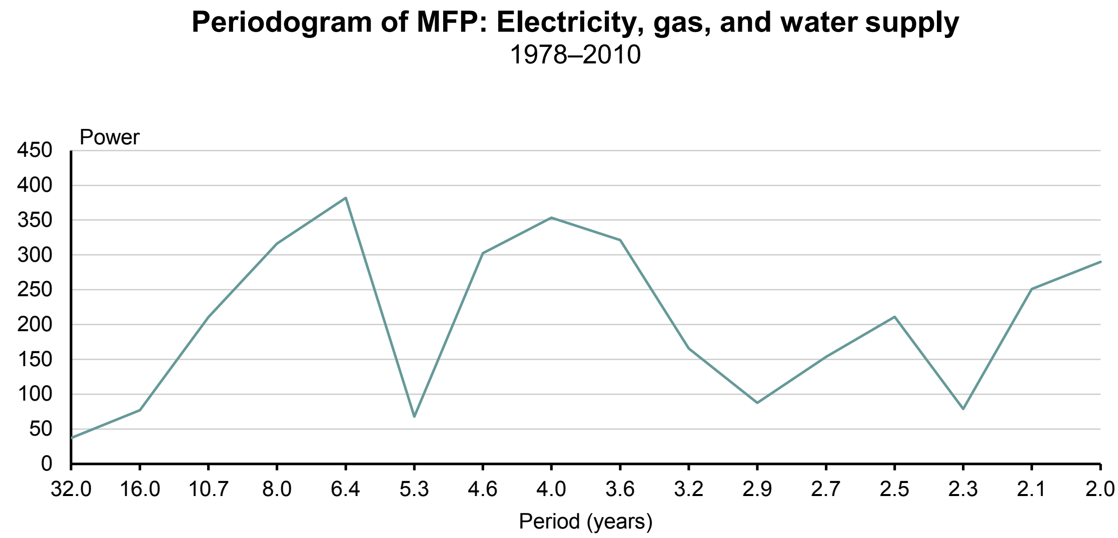

Source: Authors' calculations

The periodogram shows a power value of 382 at a period of 6.4 years and a power value of 353 at a period of 4.0 years. Unlike the quarterly GDP analysis, the two peaks here suggest two separate cycles.

We now turn our attention to providing some brief remarks on the possibility of extending Fourier analysis.

Wavelet analysis could be an extension of Fourier analysis

Our paper offers a starting point for future applied economic research using wavelet analysis. Hubbard (1998) quotes Ronald Coifman from Yale University as saying, “I view wavelet analysis as a natural extension of traditional Fourier analysis, and therefore on a scientific level a translation of mathematical tools and methods which have been in use in mathematical analysis and other sciences for the last 50 years.”

One advantage of wavelet analysis is that it is unconstrained compared with Fourier analysis. For instance, Fourier analysis works best with detrended data, but wavelet analysis works with both trending and detrended data. Wavelet analysis can also identify complicated patterns in time series, whereas Fourier analysis is unable to do so.

Wavelet analysis is used in diverse fields such as signal processing and digital imaging. But it is little used in the field of economics. Of the few promoting its use in economics, the most prominent one is James Ramsay. Ramsay (1999) investigates the properties of wavelets of relevance to the field of economics and finance. Hence our research paper can be extended by conducting wavelet analysis on the same data that we use in this paper. Doing so will confirm the periodic characteristics that we identified, and may identify other patterns that Fourier analysis is unable to detect.

Summary

We introduced our paper by stating the importance of cycles in economics. We then introduced an alternative method for studying cycles: frequency domain analysis. Frequency domain analysis has a long history, but economists have not often used it. However, the 2003 Economics Nobel Prizewinner Clive WJ Granger used it to interpret business cycles in his 1966 publication, “The typical spectral shape of an economic variable”.

We then reviewed work previously done by Statistics NZ. We provided a brief introduction to the theory underlying Fourier analysis and provided a stylised example of how to conduct Fourier analysis. We then applied this method to identify periodic behaviour in eight data sets:

- monthly retail sales
- electricity use
- labour market
- foreign currency exchange
- quarterly GDP
- labour productivity
- capital productivity
- multifactor productivity.

We provided some comments on a natural extension to Fourier analysis, namely wavelet analysis. We highlighted its advantages over Fourier analysis and provided a reference for the interested reader.

Fourier analysis can identify economic cycles

We have deliberately steered away from suggesting points in time where a cycle commenced or concluded. More precisely, we have steered away from suggesting cycle turning points. In that respect, the approach advocated by Bry and Boschan (1971) is more suitable for selecting cycle start- and end-points. Nonetheless, our approach can be used to complement theirs. Once a cycle start-point has been identified, our approach of identifying the length of a cycle can be used as an estimate for the cycle end-point. For example, Hall and McDermott (2009) identified the last New Zealand business cycle trough as the first quarter in 1998. Using our findings from the data in Hall and McDermott (2011) of a maximum cycle length of 32 quarters, we estimate the next business cycle trough to have occurred in the first quarter in 2006.

Fourier analysis can be applied to a variety of economic variables. Nerlove (1964) has used it to study the Federal Reserve Board index of industrial production and employment data. Granger and Morgenstern (1963) have used it to study stock market prices. Cunnyngham (1963) has used it to study the monthly money supply.

The strength of Fourier analysis is its ability to identify cyclic behaviour in any type of time series data. In our paper we have shown the results of Fourier analysis on a limited number of time series data to demonstrate how you can use it to detect cyclic behaviour. In reality, we conducted Fourier analysis on a large number of time series and discarded most of them. In the majority of Statistics NZ time series data we didn’t detect any cyclic behaviour at all.

References

Baxter, M, & King, RG (1999). Measuring business cycles: approximate band-pass filters for economic time series. *The Review of Economics and Statistics*, *81*(4), 575–593.

Bry, G, & Boschan, C (1971). *Cyclical analysis of time series: selected procedures and computer programs*. New York: Columbia University Press.

Chatfield, C (2004). *The analysis of time series: an introduction*. London: Chapman & Hall.

Cunnyngham, J (1963). *Spectral analysis of economic time series*. Working Paper No. 16, U.S. Department of Commerce, Bureau of Census, New York.

Granger, CWJ (1966). The typical spectral shape of an economic variable. *Econometrica*, *34*(1), 150–161.

Granger, CWJ, & Morgenstern (1964). Spectral analysis of stock market prices. *Kyklos*, *34*(1), 150–161.

Hall, VB, & McDermott, CJ (2009). The New Zealand business cycle. *Econometric Theory*, *25*: 1050–1069.

Hall, VB, & McDermott, CJ (2011). A quarterly post‐Second World War real GDP series for New Zealand. *New Zealand Economic Papers*, *45*(3), 273–298.

Hodrick, RJ, & Prescott, EC (1997). Postwar U.S. business cycles: an empirical investigation. *Journal of Money, Credit, and Banking*, *29*(1), 1–16.

Hubbard, BB (1998). *The world according to wavelets: the story of a mathematical technique in the making* (2nd ed). Natick, Massachusetts: A K Peters,.

Janacek, G (2001). *Practical time series (Arnold texts in statistics)*. London: Arnold

Nerlove, M (1964). Spectral analysis of seasonal adjustment procedures. *Econometrica*, *32*: 241–286.

Ramsay, JB (1999). The contribution of wavelets to the analysis of economic and financial data. *Philosophical Transactions of the Royal Society of London A*, *357*: 2593–2606.

Statistics New Zealand (2007). *Extracting growth cycles from productivity indexes.* Wellington: Author.