Australian National University

Crawford School of Public Policy

CAMA

Centre for Applied Macroeconomic Analysis

Turning Point and Oscillatory Cycles

CAMA Working Paper 74/2019
September 2019

Mariano Kulish
University of Sydney
Centre for Applied Macroeconomic Analysis, ANU

Adrian Pagan
University of Sydney
Centre for Applied Macroeconomic Analysis, ANU

Abstract

The early history of cycles research involved locating turning points in the data. Later, the development of methods such as spectral analysis led to a focus on oscillations. A distinction between cycles and oscillations is needed - both imply turning points, but turning points do not necessarily imply oscillations. Comin and Gertler (2006) argue that attention should be paid to medium term oscillations of 8-30 years rather than the standard 2-8 years of the business cycle, while Beaudry et al. (2019) suggest that there is a cycle of 9/10 years in series such as hours per capita. We investigate what the evidence is for the latter and find that it explains little of the variance of the data. We then show that some of the fillters being used to locate either turning points or oscillations in the series are not appropriate for the nature of the series being analyzed, specifically whether they are I(1) or I(0). Finally, we assess if the concepts of medium term and 9/10 year cycles are useful for comparing models and data. This is done by examining models of endogenous versus exogenous technology as well as limit cycles due to accumulation and complementarity mechanisms.

Keywords

Medium-term cycles, band-pass filter, business cycles

JEL Classification

C52, E32

Address for correspondence:

(E) cama.admin@anu.edu.au

ISSN 2206-0332

Turning Point and Oscillatory Cycles*

Mariano Kulish⋆ and Adrian Pagan⋄

September 2019

Abstract

The early history of cycles research involved locating turning points in the data. Later, the development of methods such as spectral analysis led to a focus on oscillations. A distinction between cycles and oscillations is needed – both imply turning points, but turning points do not necessarily imply oscillations. Comin and Gertler (2006) argue that attention should be paid to medium term oscillations of 8-30 years rather than the standard 2-8 years of the business cycle, while Beaudry et al. (2019) suggest that there is a cycle of 9/10 years in series such as hours per capita. We investigate what the evidence is for the latter and find that it explains little of the variance of the data. We then show that some of the filters being used to locate either turning points or oscillations in the series are not appropriate for the nature of the series being analyzed, specifically whether they are $I(1)$ or $I(0)$. Finally, we assess if the concepts of medium term and 9/10 year cycles are useful for comparing models and data. This is done by examining models of endogenous versus exogenous technology as well as limit cycles due to accumulation and complementarity mechanisms.

JEL Classification: C52, E32

Keywords: Medium-term cycles, band-pass filter, business cycles

*We would like to thank Dana Galizia and Tim Robinson for comments. Kulish's Research Supported by ARC Grant DP190100537. Pagan's Research Supported by ARC Grant DP160102654.

⋆School of Economics, University of Sydney. mariano.kulish@sydney.edu.au

⋄Emeritus Professor, University of Sydney and Honorary Professor, Australian National University (adrian.pagan@sydney.edu.au)

1 Introduction

For many years the seminal work of Burns and Mitchell (1946) has been an important structure upon which to develop business cycle analysis. Their work supported the idea that business cycles varied between 2 and 8 years in length and so one should focus attention upon such cycles, asking whether any model can produce cycles of that duration. In 2006, however, a dissenting opinion was expressed by Comin and Gertler (2006), who said that the focus should be on *medium term cycles* and not just the business cycle. It is not entirely clear how these should be defined. In their original paper it was a cycle between 2 and 200 quarters. Later, Comin et al. (2014) used yearly data and defined the cycle as between 2 and 50 years, distinguishing a *high frequency component* between 2 and 8 years, and a *medium term component* between 8 and 50 years. Many who have cited the first paper have used the term to refer to cycles between 8 and some number of years varying from 20 to 50, e.g Drehmann et al. (2012) use 30 years. What is common to all this work is the length of the upper bound, far greater than the 8 years of the business cycle. The argument advanced for looking at a medium term cycle was that these cycles showed very different characteristics to the business cycle, particularly in terms of volatility. More recently, Beaudry et al. (2019) have argued that one should look at what we will call a 9/10 cycle; something around 9-10 years, so longer than the 2 to 8 years of the business cycle but shorter than the medium term cycle of Comin and Gertler (2006). They argue that a 9/10 year cycle can be seen in series such as unemployment and per capita hours.

In this paper we argue that it is important to recognize that there are two ways in which cycles have been referred to in the literature - either via *turning points* in series or by *oscillations* in them. Burns and Mitchell and the NBER Business Cycle dates for recessions and expansions take the first perspective. Much academic work takes the latter. Because oscillations imply turning points in a series it is possible that the dates of peaks and troughs (turning points) reflect oscillations, but one has to be careful in drawing this conclusion, since one can have turning points in a series without a representative oscillation of interest being present.[1] Sargent (1979, p. 240) simulated data from an AR(1) process and observed that "*This illustrates how stochastic difference equations can generate processes that 'look like' they have business cycles even if their spectra do not*

[1] Because all oscillations may be present in a stationary series we need to define what we mean by a representative oscillation. Consider a white noise series. Because the spectrum is flat there is no representative (or dominant) oscillation. In the same way, an AR(1) process with a positive AR coefficient has a maximum of the spectrum at the origin. That oscillation would be of infinite periodicity and therefore not of interest to business cycle analysis. In the spectral literature a representative oscillation corresponds to where the spectrum has a maximum in a region away from the origin.

have peaks."

Section 3 then turns to how it has been proposed to find a representative cycle. Finding turning points in the series is a straightforward approach which can be applied to any series, regardless of whether they are $I(0)$ or $I(1)$. Oscillations are different. To detect them in an $I(0)$ series non-parametric methods have been used to determine if the spectrum of the series has a peak at a certain frequency.[2] They are difficult to apply when the data is $I(1)$ – Beaudry et al. (2019) suggest that this may be a reason why a 9/10 cycle would be difficult to find in GDP. Instead, they suggest it could become clearer in stationary series such as per capita hours and the rate of unemployment, as they regard these as being $I(0)$. Hence spectral methods might be applied to those series. They do so and claim to locate a 9/10 cycle. We examine this claim, as well as alternative evidence they give for a 9/10 cycle that involved finding a dominant oscillation by looking at the prediction of a recession from a past recession, where the latter are given by NBER dates.

Section 4 moves on to the question of extracting an oscillation from data. This generally involves filtering. Some of the filters used in the literature are shown to be inappropriate for the nature of the series being analyzed, specifically whether they are $I(1)$ or $I(0)$. We show that the argument that medium term cycles provides different information to that at the business cycle frequency is just an artifact of the degree of persistence in the data. Arising from the earlier spectral analysis we then ask whether, even if it exists, any 9/10 cycle in per capita hours is important? *We find that the 9/10 cycle explains only 0.9% of the volatility of hours per capita.*

Finally, section 5 looks at the use of cycle information for discriminating between models. First, we ask whether Beaudry et. al.'s criticism of DSGE models such as Christiano et al. (2014) (CMR) that they fail to replicate the 9/10 cycle in hours per capita is a reasonable one. Initially, we do this by looking at turning points in the data and from the CMR model. Then we ask whether the limit cycle models used by Beaudry et. al. do produce turning point cycles seen in the data, and whether they replicate the complete spectrum of hours, rather than just a small part of it? Even in their paper it is clear that the limit cycle model fails to replicate the spectrum and only gives a good match to a relatively minor part of it. Consequently, we question why one should only focus on a small part of it. Finally, Comin and Gertler (2006) argued that a model with

[2]There are alternatives to such non-parametric methods which involve fitting a parametric model and either using the model to determine the representative oscillation, as in say Harvey and Jaeger (1993), or providing a rule for locating turning poiints, as in MS models such as Hamilton (1989). Because these were not used by Beaudry et al. (2019) or by Comin and Gertler (2006) we do not discuss them. We also work with a single series, as that is what those authors do, although institutions such as the NBER use a range of series i.e. they find turning points in a number of series and then aggregate them into a single set of turning points. Harding and Pagan (2016) survey these multivariate approaches.

endogenous growth was superior to a standard RBC exogenous growth model when one looked at the volatility of the medium term cycle in hours, and that this superiority was not evident in the business cycle oscillations. Hence this constituted an argument for computing medium term oscillations. We find that this is not correct. As anticipated above there are some potential issues with the filter used by Comin and Gertler and this might explain the outcome. Section 6 concludes.

1.1 Finding Turning Points in Univariate Series

In much early work on cycles graphs were presented, and a visual inspection showed peaks and troughs in the series. These could be characterized by looking at local maxima and minima in the series. As Garvy (1943) notes, Kondratieff did just this. Kitchin (1923) was one of the first to be more precise about it, marking out peaks and troughs in the series he was working with, although it was never entirely clear what his rules for defining local maxima and minima were.[3] Any such definition must involve a window of time in which one sees a maximum or minimum, but it can't be too large or one will only locate the global maxima and minima in a given sample of data.

Burns and Mitchell were very clear about their identification of peaks and troughs in the levels of series.[4] They mentioned a number of factors that were used to decide on them. Bry and Boschan (1971) developed an algorithm that attempted to make a precise statement of the general principles set out in Burns and Mitchell's work, and found that their resulting turning point algorithm could do well at replicating Burns and Mitchell's decisions. A simplified version of the Bry and Boschan (BB) algorithm is the BBQ program that is an add-on in EViews, Stata, Rats, and other programs such as R.[5] It essentially follows Bry and Boschan in using a window to locate a peak and trough and then applying some censoring rules that relate to the minimal length of upturns and downturns and the complete cycle. In BB's original work, and the NBER dating of cycles, monthly data is used with a window of six months, a minimum duration for upswings and downswings of 6 months, and a minimum length for a complete cycle of 15 months. We could write that trio as 6/6/15. In quarterly terms the equivalent would be 2/2/5.[6]

[3] There is ambiguity since he sometimes over-rode the actual maxima and minima and replaced them with the peaks and troughs of what he calls the "virtual" or "ideal" cycle of 3.3 years. So his work seems to be a mixture of discovery and the imposition of a view about the length of the cycle.

[4] They always worked with series where some adjustments for seasonality and factors such as strikes were made.

[5] The simplifications used are described in Harding and Pagan (2002).

[6] For example a rule for a peak in a quarterly series y_t at time s would be $\{y_{s-2}, y_{s-1} < y_s > y_{s+1}, y_{s+2}\}$. This could be expressed as $\{0 < \Delta y_s, \Delta_2 y_s; \Delta y_{s+1}, \Delta_2 y_{s+2} < 0\}$, emphasizing it involves the signs of growth rates.

Figure 1: Quarterly Per Capita Hours

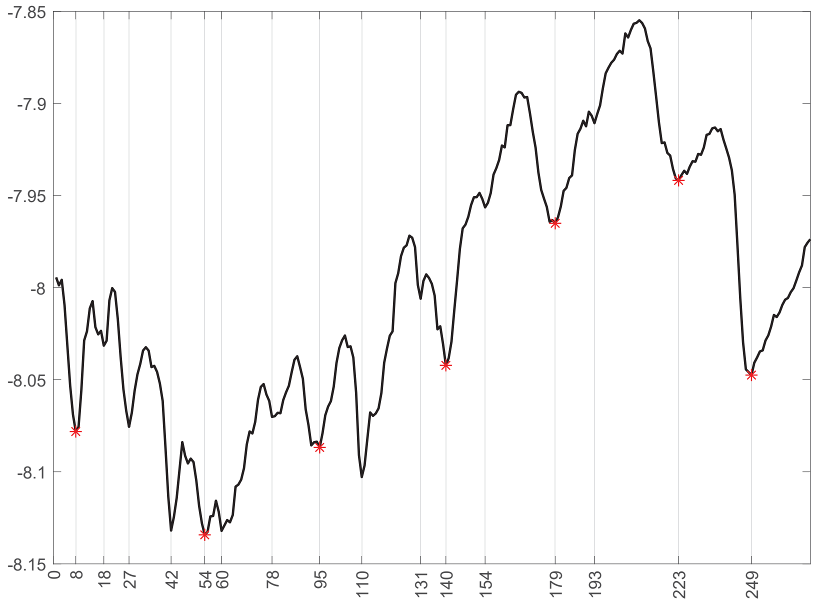

Figure 1 shows a plot of the quarterly hours per capita series used by Beaudry et al. (2019). There are clearly well defined peaks and troughs and so a well defined cycle. There seem to be 12 major troughs and then some others that are more doubtful to the eye. This data is over 1947-2017 and there are 270 quarters, thereby pointing to a cycle (distance between troughs) that is around 23 quarters long on average. How do we capture this picture with the rules referred to above? A 2/2/5 rule omits many of the minor turning points that show up as little wobbles in the graph. We can rule out others by extending it to 3/3/7. But to get anything like 36-40 quarters (9/10 years) as a cycle we need to use 7/7/15. This produces an average cycle of 40 quarters and involves troughs at observations which are denoted by the symbol $*$ in Figure 1. Looking at the figure it might be reasonable to delete those at 18, 60, 78, 131, 154 and 193, but it seems hard to justify ignoring those at 27, 42 and 110, at least based on any ocular test. So historically researchers looking at this series would have come up with a conclusion that the cycle would be between 23 and 27 quarters long.[7]

[7]Now neither the NBER nor Burns and Mitchell work with per capita quantities when dating the business cycle. The use of per capita will mean a lower growth rate in the series and a shorter cycle. The cycle we get when dating per capita GDP with 2/2/5 is 20 quarters long, so a 5 year cycle. If we use 2/2/5 on per capita hours it is 16 quarters so a 4 year cycle.

1.2 Discovering Oscillations in Univariate Series

So how did Beaudry et al. (2019) find a 9/10 year cycle in per capita hours? It clearly does not come from a study of the turning points in the data. Instead they looked for oscillations using spectral methods. They began with the periodogram which is presented in Figure 2 (we used their program). Note the big spike near the origin.[8] This is because it is a persistent series. The 10 year oscillation is highlighted in Figure 2.

Figure 2: Periodogram of Per Capita Hours

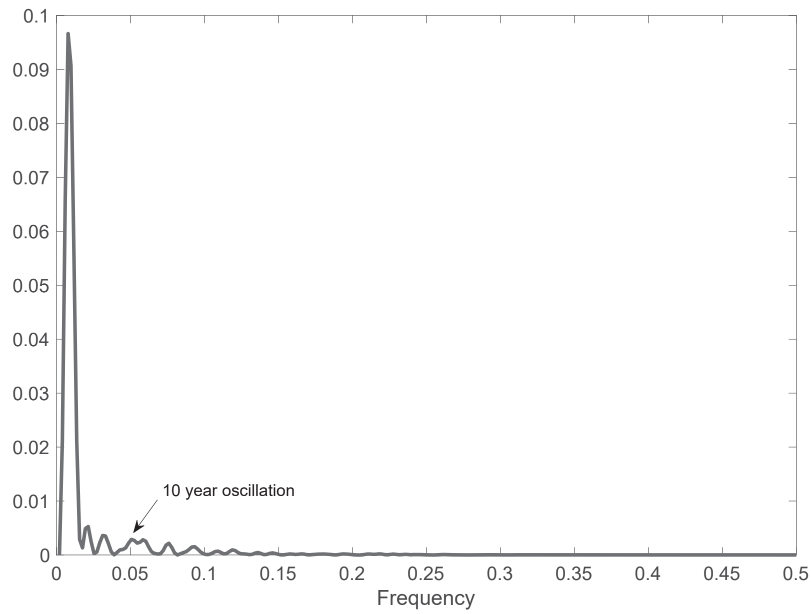

Note: for the x-axis 0.5 is equivalent to π at an angular frequency.

Now the graphs they present of spectra only involve frequencies between 1 and 15 years so the large peak near the origin is eliminated. Note that there are also other peaks before the 10 year cycle. One is at 19 years. So they cut these out by using 15 years as the upper limit in most of their graphs (their figure 11(a) labelled (2,100) being the exception). After that they smooth the periodogram ordinates using centered 13 point Hamming weights. These are symmetric and have the values 0.0870, 0.1481, 0.3152, 0.5435, 0.7717, 0.9388, 1, 0.9388, 0.7717, 0.5435, 0.3152, 0.1481, 0.0870. Figure 3 does this, resulting in the large peak near the origin being retained and the little bump at the 10 year frequency. Because the area under the spectrum is the variance of the series, it suggests that very little of the variance of hours is explained by the 9/10 year oscillation. We return to that later.

[8]It is zero at the origin since the data was mean corrected.

Figure 3: Smoothed Periodogram of Per Capita Hours,
Hamming Kernel, $M = 13$

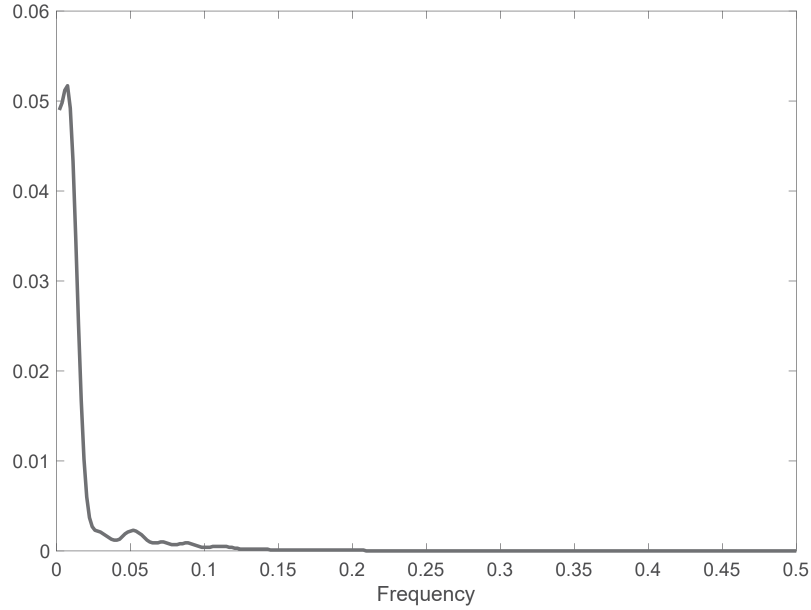

Note: for the x-axis 0.5 is equivalent to π at an angular frequency.

Now consider what the spectrum of an AR(1) with coefficient .92 looks like when computed in the same way. This is a persistent process but not one with a unit root. We simulate 270 observations to agree with their sample size.[9] Figure 4 gives this. Note that 9/10 year cycles seem to be in the data, even though they are not, and also how little power there is near the origin, even though we know that in the true spectrum of an AR(1) that is where the power is. So a 13 point averaging of the periodogram may not be enough to describe the spectrum accurately.

An alternative way of computing a spectrum is to average the autocovariances of the series. Suppose this is done with Parzen kernal weights and $M = 80$ autocovariances (they do such an analysis). Figure 5 gives this. Note the peak near the origin and a few small bumps. The first one is at 33 quarters and the next is at 7 quarters. If only 30 autocovaiances are used, Figure 6 suggests that there is no 9/10 oscillation. So the method of computing the spectrum is important for even seeing a 9/10 cycle. This is an example of the fact that, just as the turning point cycle duration depends on the precise nature of the rule selected, this is also true of spectral methods for finding oscillations. One always needs to choose a weighting scheme and the number of periodogram ordinates or autocovariances that are to be used in computing the spectrum, so it is wise to look at

[9] We follow them in padding this out to 1024 using zeroes.

Figure 4: Smoothed Periodogram for an AR(1) with $\rho = 0.92$
Hamming Kernel $M = 13$

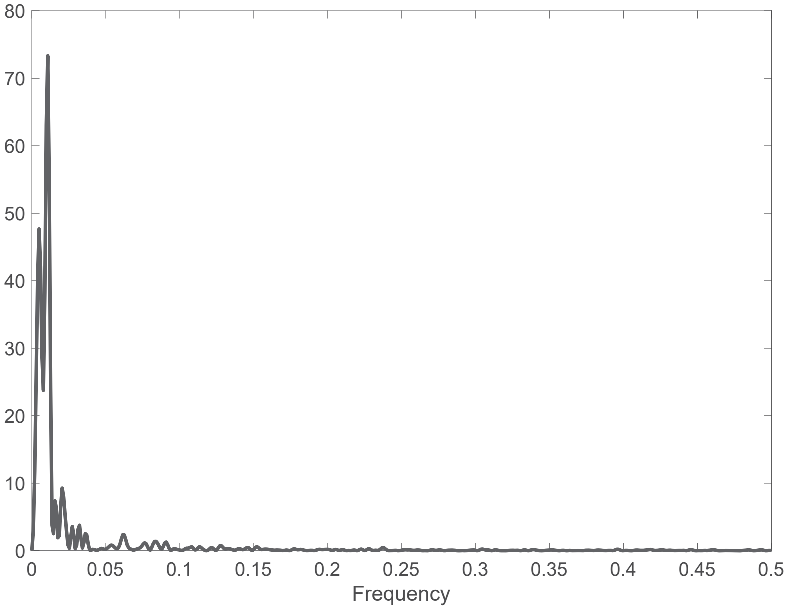

Note: for the x-axis 0.5 is equivalent to π at an angular frequency.

robustness to these choices.

1.3 Using the Probability of Recession to Suggest Cycle Length

Beaudry et. al. first suggest that there is a 9/10 cycle based on an analysis relating to the prediction of recessions into the future from a recession today. They suggest that one look at the probability of a recession at $t + k$ given that there is a recession at t. Let $R_t = 1$ if there is a recession existing at t and zero otherwise. Then that would tell us to look at $\Pr(R_{t+k} = 1|R_t = 1)$. However they do not do this. Rather the dependent variable they use is the event that at time $t + k$ there is a recession somewhere from $t + k - m$ to $t + k + m$, where $m = 3, 4$ or 5. They describe the regression with that dependent variable as a conditional probability of a recession at $t + k$ given R_t. To see the mis-interpretation here set $m = 5$ and assume that we have a sequence for R_t of {0,0,0,0,0,0,0,1,1,1,0,0,0,0,0,0,0}. Their constructed variable (called BGP_t here) from this sequence is {0,0,1,1,1,1,1,1,1,1,1,1,1,1,1,0,0}, so the BGP "recession event" is of much longer duration than is true of NBER recessions. To get the conditional probability of an NBER recession we need to set $m = 0$. Note that the BGP recession event has a sample mean that is 4 times what R is, and, if we described this as a recession event (which they

Figure 5: Spectral Density of Per Capita Hours
Parzen Kernel using 89 Autocovariances

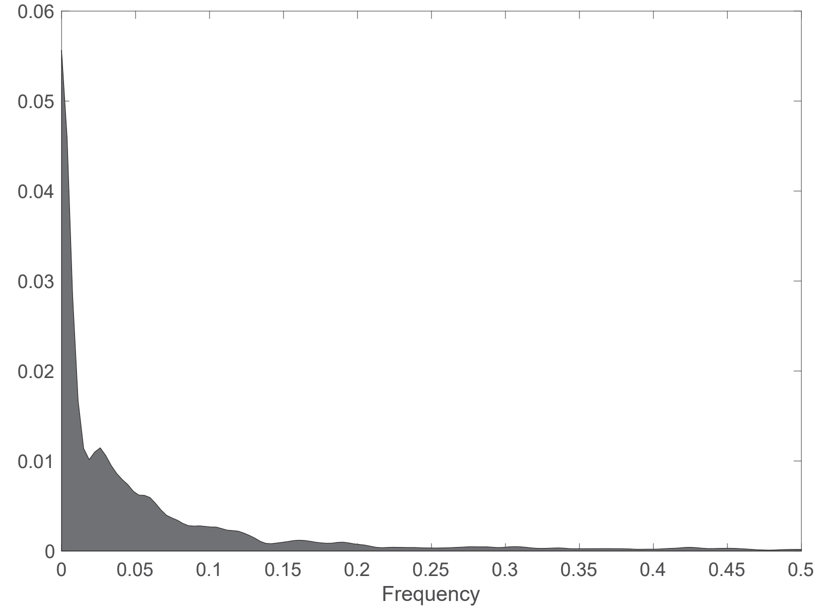

Note: for the x-axis 0.5 is equivalent to π at an angular frequency.

Figure 6: Spectral Density of Per Capita Hours
Parzen Kernel using 30 Autocovariances

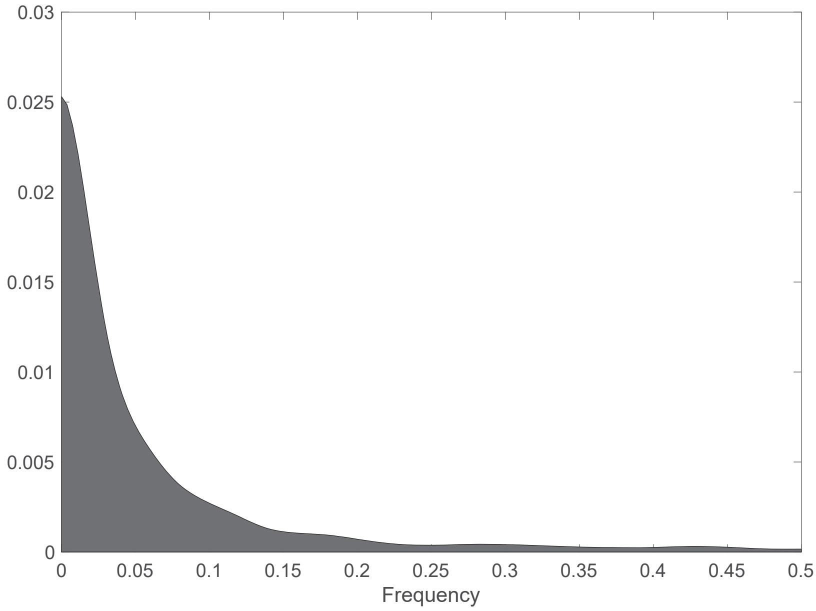

Note: for the x-axis 0.5 is equivalent to π at an angular frequency.

do), then we would be in recession most of the time.

Figure 7 gives the probability of $BGP_{t+k} = 1$ given $R_t = 1$ for $k = 12,...,59$ periods. It also gives the probability of $R_{t+k} = 1$ given $R_t = 1$.[10] Notice that the probability of the *recession event* when $k = 18$ is the maximum. This makes sense, since the average duration of a turning point cycle was 16-23 quarters. However the conditional probability is really not very high; the unconditional probability is .14. Of course for the BGP event the unconditional probability is much higher. But that is an artifact of its definition.[11]

Figure 7: Probabilities of a Recession

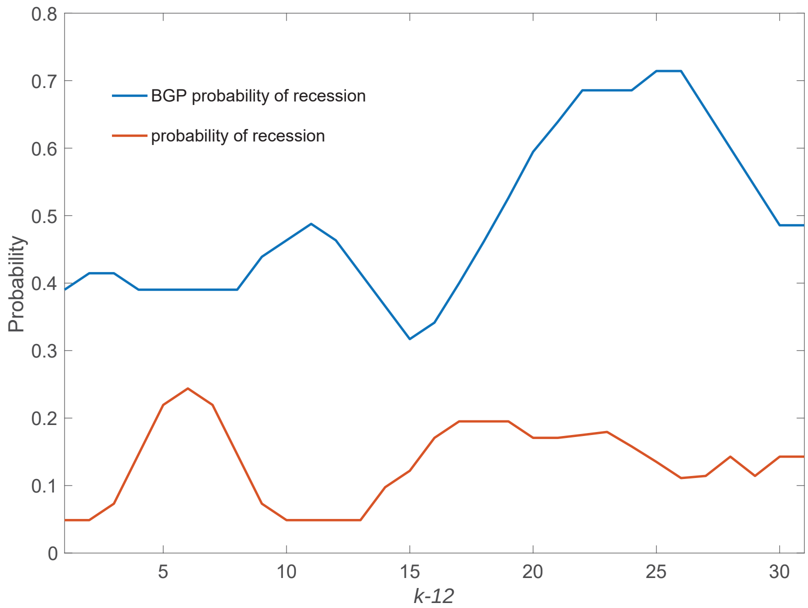

Note: $k = 12, ..., 47$.

2 Extracting an Oscillatory Component

A literature has arisen on extracting (rather than discovering) an oscillatory component. There are a variety of filters that work on the data y_t to split it into components so that $y_t = y_t^C + y_t^O$. One of these – y_t^C – is called the cycle. When y_t is $I(0)$, y_t^O would capture the remaining oscillations, given that all oscillations are present in an $I(0)$ series. If y_t is $I(1)$, then y_t^O has various names such as "trend" or "permanent". These filters are all extracted by using weighted averages of the data. The best known of these filters are

[10] The NBER dates from 1946 to 2017 for R_t were used by Beaudry et al. for this exercise.

[11] The same issue comes up in Wright (2006) who constructed BGP type events and then asked whether spreads could predict them - see section 9.4 of Harding and Pagan (2016).

from the Band-Pass class associated with Baxter and King (1999) and Christiano and Fitzgerald (2003).

A general form for the filters is[12]

$$\begin{aligned} y_t^O &= \sum_{j=0}^{T} \omega_{\pm j} y_{t\pm j} \\ &= (\sum_{j=0}^{T} \omega_{\pm j}) y_t + \sum_{j=1}^{T} \omega_{\pm j} \Delta_j y_{t\pm j}. \end{aligned}$$

Accordingly,

$$y_t^C = y_t - y_t^O = (1 - \sum_{j=0}^{T} \omega_{\pm j}) y_t - \sum_{j=1}^{T} \omega_{\pm j} \Delta_j y_{t\pm j},$$

so that, if y_t is $I(1)$ and y_t^C is to be a transitory component, then $\sum_{j=0}^{T} \omega_{\pm j} = 1$ must hold.

Christiano and Fitzgerald (CF) do provide a Band-Pass filter with this summation property (in EViews it is selected when the series is said to be $I(1)$), but many of the CF bandpass filters that are available as routines just use the weights that would be suitable only for an $I(0)$ series, and so there may be an $I(1)$ component left in what is supposed to be a transitory component when y_t is $I(1)$. To see the impact of using the wrong filter we simulate data on a pure random walk and then apply two 8-32 quarter bandpass filters of the CF variety to the series, one of which is appropriate if the series is $I(1)$ and the other if it is $I(0)$. We simulated 2000 observations and Figure 8 shows the observations from 1800 to 1960. It is clear that there are differences between the two series which gets larger as one moves towards the end of the sample. What is also particularly noticeable is that the turning points are generally different between the series. In fact, applying BBQ to the correctly filtered estimate y_t^C (that is assuming y_t is $I(1)$) we get an average oscillation of 14.6 periods, while it is only 11.9 for the incorrectly filtered one.

An application of Band-Pass filters that has been widely adopted has been Comin and Gertler (2006), who were were trying to extract "medium-term cycles". They applied the Band-Pass filter weights appropriate to an $I(0)$ series to output growth Δy_t. Then they cumulated this to account for the level of output y_t as being $I(1)$. It follows that their filtered cyclical component of Δy_t would be $\sum_{j=0}^{T} \omega_{\pm j} \Delta y_{t\pm j}$. Cumulating it we would get $\sum_{j=0}^{T} \omega_{\pm j} y_{t\pm j}$. This does not impose the restriction that $\sum_{j=0}^{T} \omega_{\pm j} = 1$, because

[12] The weights here are taken to be constant. Often they vary with time but that does not affect the argument.

Figure 8: Correct and incorrect band pass filters

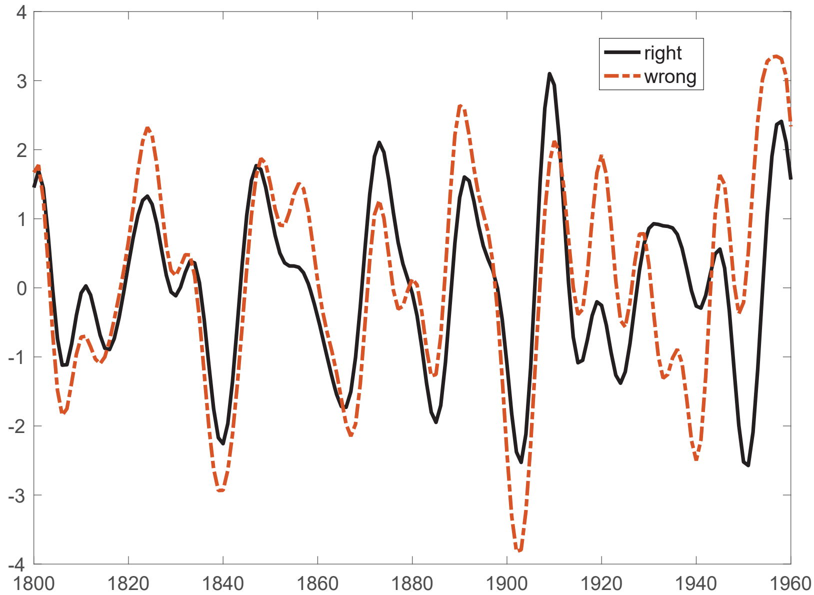

Note: 8 to 32 quarter cycles for a simulated pure random walk

the filter weights are those appropriate to an $I(0)$ and not an $I(1)$ series. Hence their extracted medium-term oscillation will include an $I(1)$ component. In contrast Beaudry et al. use the filter appropriate to an $I(1)$ series to extract medium term cycles from the per capita hours series, even though they have assumed that it is $I(0)$. So both papers utilize an incorrect filter to extract an oscillation. Figure 9 shows that the correct filtered component explains little of the per capita hours series, which we have already concluded by looking at the spectrum. Indeed it explains only 0.9% of the variance of hours.

3 What Does One Learn from Examining 9/10 and Medium Term Cycles?

3.1 From the Data

Suppose we perform a band-pass filter data while stipulating that *all* oscillations between 8 and 30 years are to be used. Call this y_t^M. This component is a mixture of all the 8-30 year oscillations. Therefore extract components y_t^j from the series to capture the oscillations $j = 8, 9, 10, ...30$ years. Then y_t^M is just the sum of the y_t^j $(j = 8, 9, ...)$. This will mean that the $8 - 30$ year cycle will tend to have turning points that are closer to those for the lowest periodicity in the range, as there are more turning points in it, and

Figure 9: Comparison of Hours per Capita and an Extracted 9/10 Oscillation

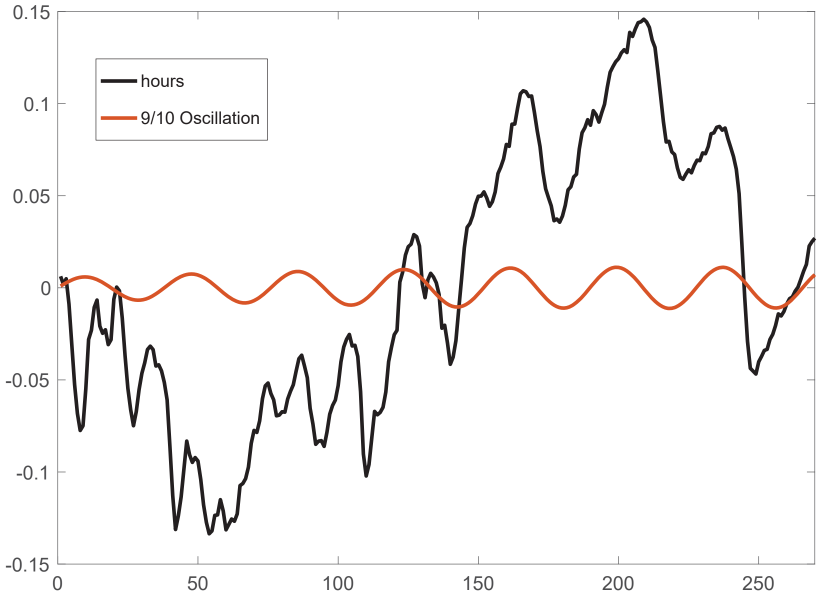

fewer at the higher end of the scale. Indeed, Comin and Gertler found that their 8-50 year cycle had a time between peaks of 10 years.[13]

What does one learn from observing (say) the component involving oscillations over 8-30 years rather than that for 2-8 years? Let us call these z_{1t} and z_{2t} respectively. A principal conclusion has been that $var(z_{1t}) > var(z_{2t})$. This led Drehmann et al. (2012) to conclude that "*...the cycles of periodicities between 8 and 30 years are more important in shaping the behaviour of these series than those with shorter duration*", while Comin and Gertler conclude that "*...the medium-term cycle in output is considerably more volatile than the conventionally measured high-frequency cycle*". To assess the contribution of a particular oscillation to the variance of a series one can examine the spectral density of the series at that oscillation. For a mixture this suggests that the area under the spectral density for oscillations between 8 and 30 years would be a measure of $var(z_{1t})$, while that between 2 and 8 years will capture $var(z_{2t})$.

To look at this more generally we simulate data on a number of processes and then compute the relative standard deviations of the 8-30 and 2-8 year oscillations. The band-pass filter used is Christiano and Fitzgerald (2003) appropriate to whether the series is $I(0)$ or $I(1)$. These results are in Table 1. It is clear that they largely reflect the amount of persistence in the data being filtered. Why does that happen? To see this consider Figure 10 which shows the spectrum for a white noise process and for an

[13] This is also the case if one dates y_t^M using the BBQ algorithm in EViews.

AR(1) process of equal variance. So the area under both curves is the same. First, take the case that the data is just white noise. Then the variance of a 2-8 year oscillation in quarterly data would be the area under the flat spectrum between the frequencies of $\frac{\pi}{4}$ and $\frac{\pi}{16}$, represented by rectangle B in Figure 10. Hence the base of the rectangle to capture the area of B is $\frac{3\pi}{16}$. In contrast, the base of rectangle A corresponding to the 8-30 year oscillation is $\frac{11\pi}{24}$ (it is between $\frac{\pi}{16}$ and $\frac{\pi}{60}$). Since the heights of the rectangles are the same, the variance of the 8-30 year oscillation relative to the 2-8 year one is 0.25, in Figure 10 A/B, and so the ratio of standard deviations should be 0.5. This is close to what we observe in the first row of results in Table 1.

Figure 10: Spectrum of an AR(1) and white noise series of equal variance

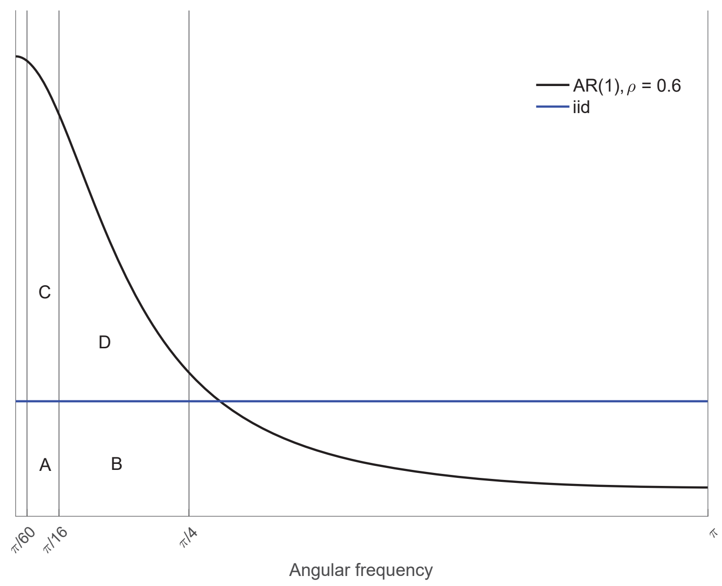

Notice that the ratio is less than unity. To make it greater than unity (as concluded by Drehmann et al. (2012) and Comin and Gertler (2006)) we need the series being filtered to depart from white noise as shown by the sprectrum of the AR(1) in Figure 10. Indeed, as Table 1 also shows, the more persistent the series the greater this ratio is. In Figure 10 this ratio in the AR(1) case is given by $(A+C)/(B+D)$. That would be expected since now there will be a greater area under the spectrum in the 8-30 year frequency $(A+C)$, as it becomes more concentrated towards the origin. Once we get to a unit root in the process we no longer have a spectrum, so the argument cannot be made precise, but the lead up to the unit root case suggests that greater persistence increases the ratio even more, which is evident in Table 1. Hence the conclusion about the relative volatilities of medium-term and more conventional business cycle frequencies is simply a consequence of how persistent the series being filtered is.

Table 1: Ratio of Standard Deviations of 8-30 year and 2-8 year Oscillations

Process	$\frac{\sigma(8,30)}{\sigma(2,8)}$
White noise	0.49
AR(1),$\rho = .6$	0.58
AR(1),$\rho = .9$	1.57
Pure random walk	2.33
Growth $\rho = .4$	2.49

Aikman et al. (2014) took UK annual data on real bank loans and GDP growth and seemed to define a medium term component using a lower limit of 8 years and a maximum of 20. They concluded that "*Consistent with the estimated density, shorter term business cycle frequency fluctuations between 2 and 8 years typically do not account for much of the overall variation in credit in the 2 to 20-year range.*"[14] If we do the same exercise as in Table 1, but now for their ratio, we get the same conclusion - as persistence rises the ratio of the 2-8 year component falls as a ratio of the 2-20 year component.

3.2 From Model Comparisons

3.2.1 Endogenous Versus Exogenous Technology Model

Extracting oscillations involves a transformation of the data and it may be that the transformed data could provide useful information when judging the quality of models. The question is whether we learn anything extra from that transformation. Comin and Gertler (2006) argued that one did. Specifically, they noted that a model which they constructed featuring endogenous technology gave a much better explanation of the volatility of hours over the whole 2-200 quarter oscillation range than RBC models did. The main exception involved hours. In connection with that they say "*The RBC model generates only about half the volatility of hours that appears in the data at either the high or medium term frequencies*" i.e. the business versus the medium term cycle. Now this comment suggests that one has *not learnt anything new* about the quality of the RBC model from the medium term analysis that wasn't evident from the higher frequency (2-32 quarters), since it seems that the volatility of hours cannot be matched there either.

[14] They supplied the data and programs in a zip file ecoj12113-sup-0002-DataS1. Some of the data had to be interpolated but this was done using their MATLAB routines. Over a range of sub-periods the mean and standard deviations for real loan growth in their table 2 match the data that we constructed.

To look at the comparison of the models further we simulated data from their endogenous technology model as well as the RBC model with a single technology shock, extracted the 2-200 quarter component (what they term the medium term cycle), and found that the RBC model *produces a higher volatility for hours* at that frequency range than their endogenous technology model. Consequently, if one is judging the quality of the models by the extent to which they can explain hours volatility, the RBC model seems better, and it is unclear why their Table 6 shows the opposite. As we noted in section 3 there are issues with the way they filtered the data, but all that is available in their replication file is the filtered data and not the original, and therefore it is hard to check why there is a discrepancy.[15] Moreover, because the RBC model doesn't produce a unit root in hours, it is unclear what the relationship between the volatility of model and data hours is, since they seem to treat hours as an $I(1)$ variable in the data.

3.2.2 Limit Cycles

Beaudry et al. (2019) have criticized existing models such as Christiano et al. (2014) (CMR) saying "... *in our exploration of different models we have not yet found an estimated model that produces a peak in the spectral density of hours similar to the one we find in the data*" (p 17). Although we have suggested that there is some doubt about a 9/10 cycle in the data we might enquire into whether the turning point cycle one would get from Christiano et al's model for per capita hours matches that in the data. Using a 2/2/5 rule we find that the length of the cycle in per capita hours from the CMR model is 13 quarters versus the 16 that we get from the actual data. Downturns and upturns for per capita hours have much the same duration as in the CMR model while, in the data, upturns are 3 quarters longer. Since CMR use normal shocks this may well explain the discrepency. Nevertheless, the CMR model simulations seem to agree quite well with the data in its cycle aspects, and it certainly does not have a 9/10 cycle, as measured by either turning points or a spectral density.

Beaudry et al. (2019) propose a model that is designed to produce limit cycles of duration 9-10 years in per capita hours. Simulating data from their model we can compute the periodogram of per capita hours. Figure 11 shows this. It is clearly quite different to that of actual per capita hours in figure 2. It is worth looking at the actual peaks and troughs in their simulated series and the 9/10 cycle extracted from it. As can be seen from Figure 12 the latter clearly misses much that is in the data.[16]

[15] The original data does not seem to be obtainable.

[16] There are other issues with the limit cycle model. It implies highly skewed simulated hours data which is not in the data.

Figure 11: Periodogram of Simulated Hours from the Limit Cycle Model

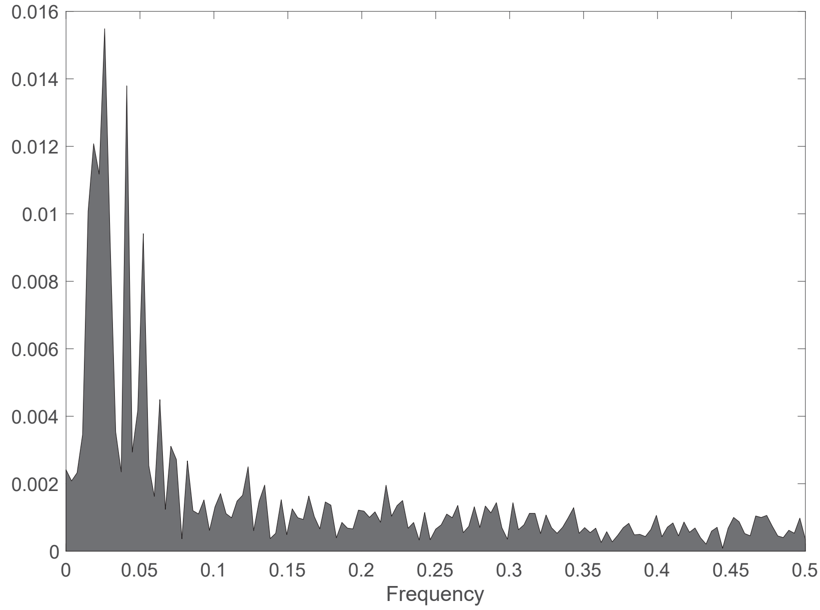

Note: for the x-axis 0.5 is equivalent to π at an angular frequency.

Figure 12: Limit Cycle Model
Simulated Hours and an Extracted 9/10 Oscillation

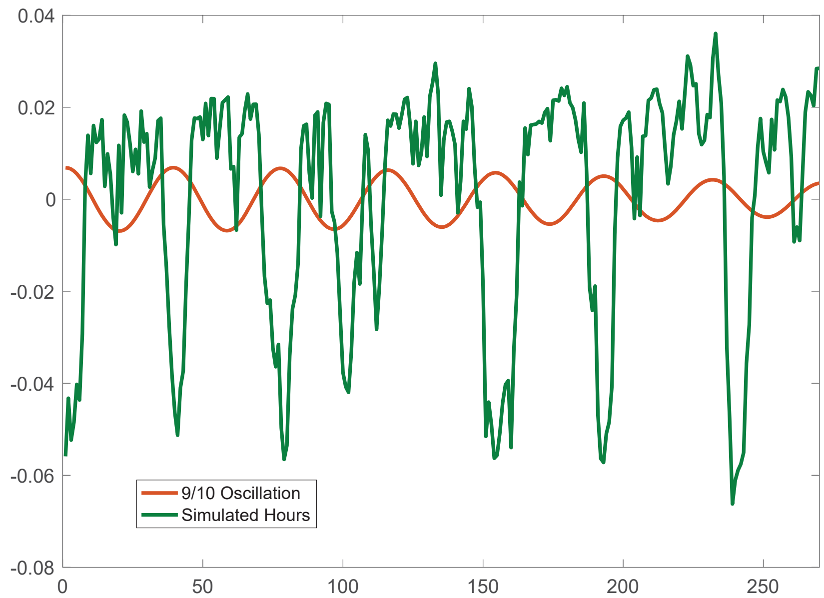

Figure 13: Spectrum of Simulated Hours from the Limit Cycle Model
Parzen Kernel, $M = 89$

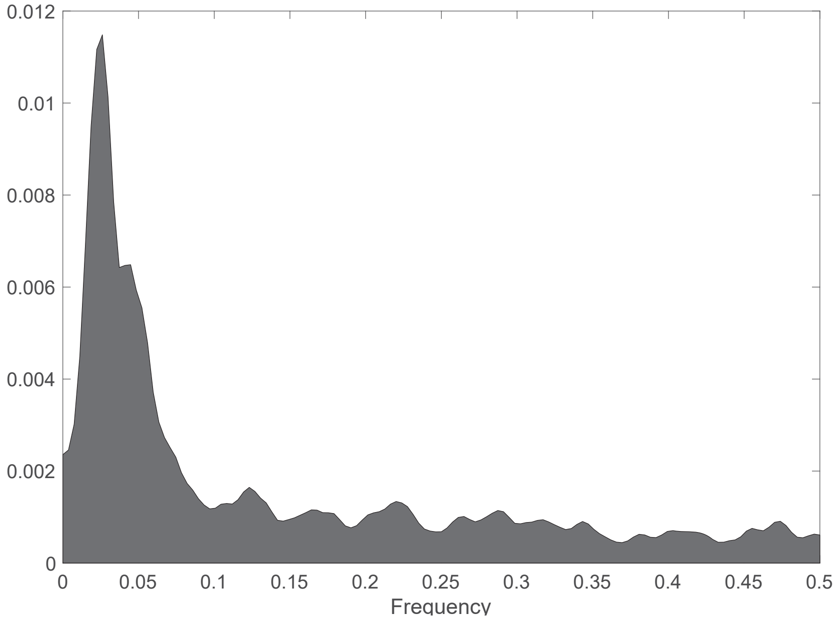

Note: for the x-axis 0.5 is equivalent to π at an angular frequency.

Figure 13 shows the spectrum that is comparable to the data in figure 5 and it is clear they are very different. It is apparent that the limit cycle spectrum matches only a small part of the data spectrum. So this raises the question of whether it is sensible to match only part of the spectrum. If we filter the per capita hours data in order to isolate a 9/10 cycle then it explains only 0.9% of the variance of per capita hours. So, unless it can also reproduce the spectrum away from that oscillation the limit cycle model is missing almost all the volatility in the series. As Figure 12 shows it can't. One possibility is that this failure comes from parameter choices for the limit cycle model. They estimated these by making a match to the spectrum over the 2-50 quarter cycle range. Would this match improve when one estimated them from a broader range of the spectral frequencies? Their figure 11 gives the model spectrum when estimation was done using 2 to 100 quarters, and it demonstrates quite clearly that there is still a major failure to capture the overall spectrum. So the question is then why should we judge a model by its ability to replicate a range of oscillations that are a minor part of the series, and not the part that has the greatest contribution to the variance? Their argument for ignoring the remainder of the spectrum is that it "...reflects largely slow-moving forces (such as demographic change) unrelated to the business cycle" If one thinks of a basic RBC model where stationary but persistent technology shocks produce a spectrum that is like an AR(1), then the spectral shape does not come from demographic change, and yet RBC models have turning point cycles like we see in the data. However, they do not have oscillations. One can make the same point about models in which there are persistent demand shocks. These will produce spectra like an AR(1) and so they are being ignored as well. It is almost as Beaudry et al believe that one cannot have a business cycle without an oscillation and the fact that this is incorrect is one of our primary points. One needs turning points in economic activity for a business cycle. These may come from a model with oscillations but need not. One cannot say that if we fail to find oscillations then we are not doing cycle research. Limit cycles are interesting for their ability to endogenously generate oscillations in series, but whether they are needed to explain actual cycles in series seems debateable based on the evidence presented here.

4 Conclusion

The paper has made the case that one needs to distinguish between two views of cycles - one involving turning points in series and one of oscillations. The former is the view used historically and which underlies the NBER cycle dates. Oscillations have turning points, but turning points can exist in a series without there being a representative oscillation of interest to business cycle analysts. One probably wants to look at both types of cycles and we have done this throughout the paper. *Business cycle analysis should not be restricted to only looking for oscillations* as often these can be of minor importance in explaining actual outcomes such as the variance of activity, as we show for per capita hours.

The use of cycle and oscillation information to judge the adequacy of models that are based more on theoretical ideas has a long history, and it can often show weaknesses not apparent from just studying a few moments of the data. Thus the recent surge of interest in looking at medium term and 9/10 oscillations rather than the business cycle might be a promising way of judging models. We did not find that this to be true. Indeed it seems to us that very little is added by taking these perspectives. It was shown that some of the cited advantages of looking at medium term cycles are really a product of easily measured characteristics of the data being filtered, in particular its persistence. We did look at the application in Comin and Gertler (2006) which argued that the medium term perspective was important, but did not find this was so. It is possible that this was a consequence of a difficulty involving the weights used to extract medium-term oscillations from $I(1)$ series. The limit cycle model developed by Beaudry et al. (2019) seems nice in theory but it does not seem to represent the data very well. It certainly produces a cycle of 9-10 years but fails at other frequencies. We have said why we doubt that there is a

9/10 cycle in the data. Even if it was then it does not seem important, explaining very little of the volatility of the series, as this depends on the complete spectrum and not a small part of it.

References

Aikman, D., Haldane, A. G., and Nelson, B. D. (2014). Curbing the Credit Cycle. *The Economic Journal*, 125(585):1072–1109.

Baxter, M. and King, R. G. (1999). Measuring Business Cycles: Approximate Band-pass Filters for Economic Time Series. *Review of Economics and Statistics*, 81(4):575–593.

Beaudry, P., Galizia, D., and Portier, F. (2019). Putting the Cycle Back into Business Cycle Analysis. *American Economic Review, forthcoming*.

Bry, G. and Boschan, C. (1971). Programmed Selection of Cyclical Turning Points. In *Cyclical Analysis of Time Series: Selected Procedures and Computer Programs*, pages 7–63. NBER.

Burns, A. F. and Mitchell, W. C. (1946). Measuring Business Cycles. *NBER Books*.

Christiano, L. J. and Fitzgerald, T. J. (2003). The Band Pass Filter. *International Economic Review*, 44(2):435–465.

Christiano, L. J., Motto, R., and Rostagno, M. (2014). Risk Shocks. *American Economic Review*, 104(1):27–65.

Comin, D. and Gertler, M. (2006). Medium-term Business Cycles. *American Economic Review*, 96(3):523–551.

Comin, D., Loayza, N., Pasha, F., and Serven, L. (2014). Medium Term Business Cycles in Developing Countries. *American Economic Journal: Macroeconomics*, 6(4):209–45.

Drehmann, M., Borio, C. E., and Tsatsaronis, K. (2012). Characterising the Financial Cycle: Don't Lose Sight of the Medium Term! *BIS Working Paper*.

Garvy, G. (1943). Kondratieff's Theory of Long Cycles. *The Review of Economic Statistics*, pages 203–220.

Hamilton, J. D. (1989). A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle. *Econometrica*, pages 357–384.

Harding, D. and Pagan, A. (2002). Dissecting the Cycle: A Methodological Investigation. *Journal of Monetary Economics*, 49(2):365–381.

Harding, D. and Pagan, A. (2016). *The econometric analysis of recurrent events in macroeconomics and finance*. Princeton University Press.

Harvey, A. C. and Jaeger, A. (1993). Detrending, Stylized Facts and the Business Cycle. *Journal of Applied Econometrics*, 8(3):231–247.

Kitchin, J. (1923). Cycles and Trends in Economic Factors. *The Review of Economic Statistics*, pages 10–16.

Sargent, T. J. (1979). *Macroeconomic theory.* Academic Press, New York.

Wright, J. H. (2006). The Yield Curve and Predicting Recessions. *Finance and Economics Discussion Paper*, 2006-7, Board of Governors of the Federal Reserve.