



Detrending and business cycle facts: A user's guide

Fabio Canova^{a,b,c,*}

^a *Department of Economics, Universitat Pompeu Fabra, c/Triass Fargas 25–27, 08005 Barcelona, Spain*

^b *Department of Economics, Università di Modena, 41100 Modena, Italy*

^c *CEPR, London, UK*

Received 17 October 1997; accepted 17 November 1997

Abstract

This note argues that it is hardly the case that the profession is fully aware that the application of different filters to the data leads to different outcomes and that we have enough evidence to claim that these differences are unimportant to evaluate the fit models to the data. © 1998 Elsevier Science B.V. All rights reserved.

JEL classification: B41; E32

Keywords: Business cycles; Detrending; Model evaluation

Stories happen to those who are able to tell them.

Paul Auster

1. A brief history of the crusade: on the generics of Burnside's criticism

When I started writing *Detrending and Business Cycle Facts* I could not forecast that I was opening a can of worms and that it was going to be so difficult to get, what I thought was a simple message, through to the profession. Now seven plus years and several referees later, I find myself in need to explain the genesis (and the catharsis) of the paper, to try to rephrase one more time what the paper is all about, and, perhaps most importantly, to clarify what the paper does not do. Thinking about what has happened since the beginning of this enterprise, I believe that I have a good story to tell at cocktail parties of the next millenium.

* E-mail: canova@upf.es

Several years ago, I was asked to be the discussant of one of the papers of Backus, Kehoe and Kydland at the Winter Meetings of the Econometric Society. I was impressed by the purity and the formal beauty of their science – get facts from the data, build fully articulated economies which give us an understanding of what is the mechanism which allows us to match or not to match them. The result was that I did not have much to say about the paper. At that time HP filtering of the data was ‘the’ approach used by business cycle theorists. Since I was not too sure about what the filter was doing in practice, I started playing around with other filters, popular in other branches of applied macroeconomics, filters whose producers or users claimed, in one way or another, were extracting reasonable cyclical components of the data.

One of the participants of that session, after hearing the discussion which was the embryo of the paper, told me ‘You are really a skunk in a rose garden’ and since then that foul smell followed the paper around. Comments both from commentators at conferences and referees, have typically been sanguine (to put it mildly) and lined up on two opposite and equally defensive positions. There were those who claimed I had discovered hot water – we all knew about it, why should we be surprised at all? And those armed with intellectual skepticism who thought I was wasting their time – for the things we care, it really does not matter.

Burnside seems to belong to the restricted group of people who effectively synthesize the two types of criticisms to the paper (we all knew about it and anyway, once you do things properly, it does not matter) and some of his arguments may have some merit to the extent that the most sophisticated portion of the profession has indeed substantially refined both the empirical analysis (we now see pictures of spectra and coherences instead of tables of variances and correlations, e.g. King and Watson (1996)) and the way the matching between the model and the data is done.

However, one of the reasons for his adverse opinion of the paper rests on the fact that ‘Economists will be misled only to the extent that they believe that all filters designed to extract the ‘cyclical’ and the ‘trend’ components of time series produce the same outcomes’. I would argue, contrary to Burnside’s belief, that this is not an unusual circumstance. In fact, it is still very common to find published papers in major academic journals with sentences like ‘stylized facts are broadly insensitive to detrending’ or ‘filtering the data with FOD, HP or linear filters produce similar results’. One may interpret this as a signal of carelessness on the part of researchers but I am inclined to think that there is still a deeply rooted credence that stylized facts do not depend on the way the data is filtered.

Furthermore, Burnside writes ‘... a theory in which productivity and hours ‘comove’ may or may not be consistent with (both of) these facts, because the statement ‘comove’ is simply too imprecise to be of use’. But this is not the statement of a typical applied business cycle analyst. Such a person is more likely to believe that there is *one set of facts* (hours is less volatile than output,

productivity is uncorrelated with hours, etc.) to be explained by theory. Burnside is also doubtful about the task set up in the introduction. He finds it hopeless, given the nature of various filters, to try to find a set of facts which is invariant to the definition of cycle employed. If a set of relationships is approximately constant, say, at business cycle frequencies, then most filters should give more or less the same result and this should deemphasize the issue of what filter we should use in documenting business cycle facts. Also, we know that models are unlikely to be able to generate coherences and phase shifts which widely vary across a narrow band of frequencies. As hopeless as it could be, by concentrating on relationships which are invariant over these frequencies, we avoid the proliferation of facts to be explained (e.g. the coherence and the phase for cycles of 6–8 yr have these characteristics, for cycles of 4–6 yr they have these other characteristics, etc.) which makes the task of theorists almost impossible and weakens the relevance of the methodological approach.

2. Should I say it or should I not? What the paper does

The paper tries to give an honest and comprehensive account of the possible outcomes of filtering on a standard set of macroeconomic series, taking an agnostic (flat prior) point of view about the DGP of the actual data. It also attempts to explain why we have seen a large number of papers reporting contrasting evidence on the same issues, even when the same data set is used. Finally, by spelling out in detail the assumptions underlying each decomposition, it allows us to go beyond simple documentation of the differences and interpret them in terms of differences in the resulting spectral decomposition.

The audience to whom it is directed consists of applied macroeconomists (i) who were not fully aware of the fact that filtering massages the data in, sometimes, unpredictable ways; (ii) who think that ‘theory is crucial in selecting which facts to report’ (see Kydland and Prescott (1990) (p. 3)) and, as a consequence, that facts are simple statistics, more or less self-evident, which display interesting and robust patterns; and (iii) who believe that one filtering method is good for all purposes. Sophisticated practitioners who are fully conscientious of the effects of filtering, of the multiplicity of possible outcomes and of the danger of the exclusive use of one filter were not the target of this paper and were always kindly excused from stepping out during the presentations.

Moreover, as is stated in the introduction, the focus of the paper is narrow and the product limited. In particular, the paper is not expected to revolutionize the status of applied macroeconomics, but simply warns users of standard techniques of the potential problems they face.

More importantly, the paper is not concerned with the evaluation of models, as Burnside seems to believe, but only with the preliminary step of collecting a set of interesting statistics of the actual data. I thought that this was the best

way to focus attention on this important problem, leaving aside other issues which are further downstream (e.g. formal evaluation of models). By doing so the paper may turn out to be vulnerable to one of the criticisms that Burnside makes: there are formal evaluation methodologies which can disregard the results of the paper. However, as I will argue later on, this is still to be demonstrated. Also, if someone believes that collecting facts is important and valuable *per se*, regardless of whether it is used at a second stage to evaluate a model, then the exercise forces researchers to go beyond simply pro or counter cyclicity to specify the frequencies where these phenomena occur (as Burnside does) and to attempt to focus, at least as a first approximation, on facts which are robust across frequencies of interest. Finally, the paper does not provide an alternative empirical methodology, but rather a few simple hints on how either to minimize the distorting effects of filters in reporting statistics of the data or to avoid inappropriate statements based on the limited evidence provided by one or two all-purpose filters.

3. Are we on opposite sides? On the details of Burnside's criticisms

In a nutshell, Burnside's criticisms of the paper is based on three arguments:

- There are several regularities in the data.
- The sign and magnitude of the moments are sensitive to detrending, but we should not be concerned with this variety as long as we are aware of the differences induced by alternative filtering methods.
- Applied macroeconomists should be concerned with whether these differences affects the testing of models. In particular, they should care to know whether certain approaches lack power in distinguishing alternatives. Because the formal statistical methodology he proposed is clearly able to distinguish alternatives, we should be even less worried about the sensitivity of business cycle statistics to detrending.

In the rest of this section I will argue that, while I generically agree on the first two points, provided that the issues at stake are well displayed on the table, the formal statistical approach proposed by Burnside is unlikely to solve the crucial problem of documenting regular features of business cycles. In addition, it does not shed light on the issue of what is the effect of filtering the data with different procedures and suffers from methodological weaknesses which, at best, make me doubt its usefulness to evaluate models against the data.

It is clearly not in the interest of anyone to be an anarchist in the field (there are no facts which are robust; tell me what fact you want and I will find a detrending method which will give it to you), but it is also true that ignoring the differences that do exist only leads us to build a falsely secure or consolatory

version of quantitative macroeconomics. Once we have accepted this step, we could quibble on whether the glass is really half full or half empty – are the second order properties at business cycle frequencies of the four data series used by Burnside robust? I would say that those for consumption and investment are broadly robust, while those for hours and average productivity are not. Whatever position one takes, I strongly believe that the more systematic evidence we have the better. On this ground, the documentation that the paper provides is the first comprehensive attempt of this type. I do not think Burnside denies this point nor do I believe he is against reporting surprising or unpleasant outcomes. The paper emphasizes the non-robust aspects of the data so strongly because they went against my prior belief. I thought, as many others did, that there was not much variation at business cycle frequencies. Hence, I expected to find much more uniformity in the results, at least across procedures which extract cycles of similar length.

Second, as I have already pointed out, I believe that applied macroeconomists are largely unaware of the differences generated by different filtering methods. Clearly, if the class of people suffering this myopia has measure zero, the exercise I conduct is irrelevant. To the extent that this is not the case, stressing that different filters give different outcomes it is not a lost battle in two senses: first, by pointing out that differences are not minor it forces researchers to take a stand on a number of assumptions which are always in the background of applied work but never mentioned (e.g. what is the theoretical correlation between the cyclical and the trend components). Second, it should raise concern among those researchers who either mechanically use standard filters, at times throwing away the baby with the water (see also Ravn and Uhlig (1997) on a related point), or believe that facts are constant within business cycle frequencies.

I would also be very surprised if more than a few researchers appreciate Burnside's fine distinction between detrending (eliminating the trend from the series) and filtering (extracting the cyclical component of the series) and the notion that the two concepts are operationally different. Along the same line, is it really true that an econometrician who uses the FOD filter to render a time series stationary is unlikely to argue that, what he has recovered after the transformation, is a measure of the cyclical component of the series? There are many examples in the literature of ARIMA based detrending filters, which approximately eliminate the random walk component of the data, whose residuals are taken as an estimate of the cyclical component (see, e.g. Beveridge and Nelson (1981), Watson (1986), Evans and Reichlin (1994) or even Hamilton (1989)). Chapter 1 of the textbook of Blanchard and Fisher (1989) is another clear example where residuals of univariate ARIMA processes are taken to represent the cyclical components. Hence, if the difference between filtering and detrending was truly well understood in the profession, part of this paper would be useless. To the extent that it has increased the awareness of the researchers of this problem, I consider the paper worthwhile.

Third, Burnside proposes an applied methodology (see Fig. 6), that is contrary to what he believes I suggested in the conclusions of the paper. If followed, his methodology makes the diversity documented and the conclusions of the paper completely irrelevant. I sympathize with the general philosophy of the approach he proposes and I do not see that my suggestion for a more interactive relationship between theory and practice is inconsistent with this philosophy. Certainly I welcome empirical exercises where business cycle tests are used to formally distinguish alternative theories of the business cycle (see Canova, 1995). However, I disagree on the usefulness of the approach outlined in Section 3 of Burnside's reply to (i) document features of business cycles, (ii) assess the likely effects of different filters in testing theories, (iii) operationally provide an evaluation procedure which allows us to distinguish the conformity of theories to the data.

Point (i) is self-evident and does not require much discussion. If one cares about documenting the statistical features of actual business cycles, the methodology Burnside outlines does not help researchers engaged in this activity. On the other hand, systematically discerning robust and non-robust facts may help researchers to organize the presentation of results for public consumption. Points (ii) and (iii) are directly linked to an underlying and contradictory stated assumption about what an economic model is. Burnside claims at one point '*... that most models are presumed to be false*' but proceeds as if a model were the correct description of the DGP of the data. I believe that this confusion is the root of the problem with the suggested approach. Let me briefly summarize the main features of his exercise. He takes two models and does the following:

- Take model $i = A, B$ as the DGP of the actual data, simulate data, filter them with HP and FOD filters and compute a set of statistics.
- Simulate data from models A and B , after the relevant parameters have been estimated with just identified GMM, filter them with HP and FOD filters and compute the same set of statistics.
- Repeat the previous step 1000 times.
- Calculate, for each DGP, the percentage of rejections of a Wald test examining whether the statistics produced by model $i = A, B$ are the same as the statistics of the actual DGP.

In other words, Burnside provides us with a nice Monte Carlo example demonstrating the size and the power of standard econometric tests when one of the two models is the true DGP of the data. In particular, the exercise shows us that estimating the parameters of the model has no influence on filtered statistics. This evidence would be relevant for the issues discussed in the paper if the crucial underlying assumption of correctness of one of the two models is appropriate. But, do we really believe that an economic model is the correct DGP of the actual data? I would say no. Then, what would happen if both

models that Burnside uses are a misspecified description of the data? The exercise does not give us any hint of what is the power of the proposed test in this situation. Moreover, the exercise sheds no light on what would happen when data from different misspecified models is filtered with alternative filters because it is not hard to think of situations where some models may be more misspecified than others at certain frequencies, either because the transmission mechanism of shocks is too rudimental or because certain models are built with an eye at specific (stylized) facts previously spotted in the data. This may be an area of interesting research for the future, but the tools used by Burnside would not help us make progress in this direction. Wald tests are inappropriate if there is no available model which may have generated the data (roughly, there is an unknown non-centrality parameter which biases testing towards rejection). My guess is that Wald tests are even more inappropriate when filtered data from misspecified models are used in constructing the tests. Note that this problem is entirely independent of how parameters are estimated (the additional problems which may arise with GMM estimation are well known to Burnside and all of those who have contributed to increase our understanding of the mechanics of this estimation approach, see, e.g. special issue of the *Journal of Business and Economic Statistics* of July 1996 devoted to this question).

To conclude, the Monte Carlo exercise conducted by Burnside has very little bearing with the two main questions that applied business cycle researchers are concerned with, document features of actual business cycle and evaluate the fit of misspecified models to the filtered data. Having said so, I have to admit that I do not have a fully articulated and comprehensive alternative able to address both issues. The silence of the paper, especially on the second problem, is partly a reflection of this ignorance. I believe that the idea of analyzing the effect of filtering the data and of systematically recording the resulting differences, together with some of the testing ideas I have suggested in Canova (1994) and (1995), constitute the skeleton for an interactive empirical approach which is more sensible than the one suggested by Burnside.

Acknowledgements

I would like to thank Morten Ravn, Russell Cooper and Adrian Pagan for endless conversations on the issues discussed in the paper.

References

- Beveridge, S., Nelson, C., 1981. A new approach to decomposition of economic time series into permanent and transitory components with particular attention to measurement of the business cycle. *Journal of Monetary Economics* 7, 151–174.

- Blanchard, O., Fisher, S., 1989. *Lectures on Macroeconomics*. MIT Press, Cambridge, MA.
- Canova, F., 1994. Does detrending matter for the determination of the reference cycle and for the selection of turning points? *Economic Journal*, forthcoming.
- Canova, F., 1995. Sensitivity analysis and model evaluation in dynamic general equilibrium economies. *International Economic Review* 36, 477–501.
- Evans, P., Reichlin, L., 1994. Information, forecast and the measurement of the business cycle. *Journal of Monetary Economics* 33, 233–254.
- Hamilton, J., 1989. A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica* 57, 357–384.
- Journal of Business and Economic Statistics* (1996) Special Issue on GMM Estimation, July.
- King, R., Watson, M., 1996. Money, prices, interest rates and the business cycle. *Review of Economics and Statistics* LXXVIII, 35–54.
- Kydland, F., Prescott, E., 1990. Business cycles: real facts and monetary myth. Federal Reserve Bank of Minneapolis, *Quarterly Review*, Spring 3–18.
- Ravn, M., Uhlig, H., 1997. On adjusting the HP filter for the frequency of the observations. Tilburg University Working Paper.
- Watson, M., 1986. Univariate detrending methods with stochastic trends. *Journal of Monetary Economics* 18, 49–75.