

INSTITUT DE STATISTIQUE
BIOSTATISTIQUE ET
SCIENCES ACTUARIELLES
(ISBA)

UNIVERSITÉ CATHOLIQUE DE LOUVAIN



DISCUSSION
PAPER

2011/32

**ASYMMETRIES IN BUSINESS CYCLES AND
THE ROLE OF OIL PRODUCTION**

DANIEL, B.C., HAFNER , C.M., MANNER, H. and L. SIMAR

Asymmetries in Business Cycles and the Role of Oil Production

Betty C. Daniel^{*1}, Christian M. Hafner², Hans Manner³, and Léopold Simar²

¹*Department of Economics, The University at Albany*

²*Institut de statistique and CORE, Université catholique de Louvain*

³*Department of Social and Economic Statistics, University of Cologne*

September 14, 2011

Abstract

Innovations to total factor productivity are thought to be an important determinant of business cycles. We investigate whether innovations to quarterly HP-filtered Solow residuals are symmetric over time for eleven OECD countries. Our model modifies classical stochastic frontier analysis to accommodate the strong serial correlation in macro data. The results have implications for whether business cycles are symmetric, with the economy responding in a linear way to normal iid shocks, or asymmetric with recessions fundamentally different from booms. Likelihood ratio tests imply that nine of eleven countries have significant asymmetries. We also consider structural differences in economies with and without asymmetries and find that asymmetries tend to be stronger the less oil production per worker in the economy. Non-linear conditioning of the HP-filtered Solow residual on the relative price of oil removes or reduces asymmetries for most countries which otherwise exhibit them, implying that much asymmetry is due to the response of the economy to oil prices.

Keywords: Solow residuals, stochastic frontier analysis, oil prices, business cycles

JEL Classification: E32, C22, C13

*Corresponding author, b.daniel@albany.edu

1 Introduction

Are business cycles symmetric? Since shocks to total factor productivity (TFP), measured as Solow residuals, are thought to be a major determinant of business cycles¹, we can narrow the question by asking whether shocks to Solow residuals are symmetric. If positive and negative shocks have the same fundamental cause, then symmetry is a reasonable presumption. Early real business cycle models (RBC) and many later dynamic stochastic general equilibrium models (DSGE) identify shocks to total factor productivity as technology shocks. The presumption of symmetry leads to linearized dynamic models, with booms caused by positive technological innovations and recessions caused by negative ones. By implication, the worldwide recession and financial crisis, which began in 2008, was caused by a very large negative technological innovation. Economists are likely to work for decades to understand the source of this large negative innovation, as they have in seeking to understand the source of the large negative innovation causing the Great Depression. Technological regress seems fundamentally different from technological progress. Perhaps negative TFP shocks are due to different fundamental factors with different distributions, requiring that we take seriously the possibility of asymmetries in business cycles.

Greenwood et al. (2000) spawned a large literature arguing that Solow residuals confound sectoral productivity shocks, and that productivity shocks in investment industries have different dynamic effects from those in consumption industries. This literature argues that disaggregation of productivity shocks is essential to understanding business cycles, but retains the presumption that technology shocks are the source of the productivity shock and are the fundamental cause of business cycles.

There are two primary problems with representing Solow residuals as technology shocks. First, a negative Solow residual shock, which is larger than the trend, is a negative shock to technology, and we do not understand technological regress.² Second, the benchmark neoclassical model has no place for excess capacity, which could be created by labor hoarding, variable labor effort, or less than fully-utilized capital stock. Since measurement of the Solow residual does not adjust for capacity utilization, Solow residuals could combine true technology change with adjustments to capacity utilization.

These problems generated a large literature in the mid-1990's, attempting to under-

¹Even if they are not the exogenous determinant, Solow residuals are highly correlated with output.

²Bad weather and increased regulation are technological regress, but only rarely do we attribute particular recessions to these causes.

stand what generates TFP shocks, particularly negative ones. The dominant hypothesis that emerged from this literature was that measured Solow residuals are some combination of true technology shocks and the response of endogenous capacity utilization to technology and other shocks in the economy. This was justified on theoretical and empirical grounds. Many authors demonstrated that numerous alternative real-world complications to the benchmark neoclassical model, including costs of adjusting factors of production (Bils and Cho 1994), capital depreciation dependent upon utilization (Burnside and Eichenbaum 1996), and costs of commuting and endogenous work effort (Burnside et al. 1993, Bils and Cho 1994, Burnside and Eichenbaum 1996), could imply optimal endogenous adjustment of capacity utilization in response to shocks. When a firm experiences a fall in demand, it can choose to reduce capacity utilization. The same firm will find it optimal to operate with some excess capacity on average so that it can respond quickly to positive shocks, implying that adjustments in capacity utilization in response to small shocks should be symmetric. Solutions for these models usually require linearization, restricting their applicability to small shocks. However, if shocks are large, the possibility for asymmetry arises since capacity constraints could bind in the upward direction, but not in the downward direction.

The hypothesis that Solow residuals confound changes in capacity utilization with true technological change receives empirical support. Evans (1992) demonstrates that demand shocks predict Solow-residual shocks. Basu (1996) uses materials inputs as a measure of capacity utilization and shows that measured Solow residuals are inversely related to materials inputs, implying that an increase in the Solow residual measures a reduction in capacity utilization

An alternative hypothesis in the literature as to the cause of negative productivity shocks is that they are due to an increase in the price of imported oil. Hamilton (1983, 2011) demonstrated that all but one post World War II recession in the US were preceded by a large rise in oil prices. Since recessions also tend to be accompanied by negative Solow residuals, then oil prices and Solow residuals should have a negative relationship. An increase in the price of oil for an oil-importing country deteriorates its terms of trade. Kehoe and Ruhl (2008) document the close negative empirical relationship between productivity, measured as Solow residuals, and the terms of trade. Additionally, imported oil is an intermediate input in the production function, and when the country uses less oil due to an increase in its relative price, production, relative to capital and labor, falls. However, they argue that since the computation for GDP adjusts for imported interme-

mediate inputs, the reduction in production, due to less use of imported oil, is exactly offset by a reduction in imported intermediate inputs, implying that there should be no fall in GDP, for given quantities of labor and capital, and therefore no fall in the Solow residual. Kehoe and Ruhl's model abstracts from variable capacity utilization. Perhaps the deterioration in the terms of trade, which reduces the purchasing power of domestic residents over world goods, reduces the demand for domestic goods, endogenously reducing capacity utilization instead. Small oil price shocks should have symmetric capacity utilization effects, but large ones might not.

Hamilton has long advocated both for asymmetries in business cycles, using Markov-switching models (Hamilton 1989, 1994), and for the importance of large oil price increases causing recessions (Hamilton 1983, 2011), while large oil price reductions do not cause similarly sized booms. His oil price models suggest that the response of the economy to oil price increases generates the asymmetry, not that the shocks themselves are asymmetric, whereas the Markov-switching models allow the possibility of different distributions of shocks in booms and recessions. His argument for oil prices is empirical, but could be theoretically explained by models which allow endogenous changes in capacity utilization in response to changes in demand conditions, created by terms of trade changes. Capacity can fall indefinitely in response to a large oil price increase, but the immediate increase in capacity utilization is bounded.

In this paper we determine whether Solow residuals, a major source of business cycles, have an asymmetric negative component. Asymmetries would imply that the Solow residual measures something other than symmetric technology shocks, possibly asymmetric responses of capacity utilization to large shocks. We introduce stochastic frontier analysis, borrowed from the micro literature on productivity analysis, as a method for decomposing innovations to Solow residuals into symmetric and an asymmetric components. This technique assumes that the innovation is a composite of a one-sided negative error and a normal two-sided error, and it is possible to estimate the ratio of variances for the two components of the innovations. If Solow residuals are symmetric, then we should attribute no variance to the one-sided error. We compare the restricted model without asymmetries to the unrestricted one using a likelihood ratio test.

To apply the stochastic frontier methodology to aggregate time series data, we must amend the micro technique to deal with both the trend and the strong persistence in macro data. Following the RBC and DSGE literature, we detrend Solow residuals using the HP filter to yield stationary series. We modify the standard stochastic frontier

model to allow for autocorrelated composite errors. After experimenting with alternative representations for persistence, we chose to treat the detrended Solow residuals as an autoregressive process with a composite error term, thereby leaving the source of the persistence unspecified.

Our sample consists of eleven OECD countries which had quarterly data available back at least to the early 1980's on GDP, employment, and investment. We find that for nine of the eleven countries, the model with an asymmetric error outperforms the model with a symmetric error. Therefore, our analysis shows that measured Solow residuals are asymmetric for most countries, implying that understanding business cycles requires models which allow asymmetries. Additionally, we find that the extent of asymmetry in the innovations to Solow residuals varies considerably by country.

Next we try to determine why asymmetry varies across countries. We observe that asymmetries are relatively larger for the group of countries with little oil production per worker. This leads us to reestimate the model after non-linear conditioning on oil prices to allow threshold effects, similar to Hamilton (2011). We find that for most countries which have asymmetric Solow residuals, oil prices are significant, either as a linear term or as a non-linear threshold term measuring large oil price increases. For most of these countries, conditioning on oil prices eliminates or reduces the asymmetry in the Solow residual. The results for these countries support Hamilton's hypothesis that recessions are caused by large oil price increases, while similarly-sized oil price decreases do not cause symmetric booms. Hamilton's hypothesis is not supported for countries which produce large amounts of oil per worker.

The rest of the paper is organized as follow. In Section 2 we introduce our econometric methodology. Section 3 presents the data and empirical results, and Section 4 concludes.

2 Methodology

In this section, we develop an econometric model to estimate whether Solow residuals have an asymmetric component. We modify the classical parametric stochastic frontier model, which was simultaneously proposed by Meeusen and van den Broek (1977) and Aigner et al. (1977), to estimate inefficiency across individual firms. Modification is necessary because data on Solow residuals have a trend and significant autocorrelation, unlike cross-sectional firm data.

2.1 Model specification

We assume that a time series of Solow residuals, given by A_t , has three components according to

$$A_t = \exp(s_t + \mu_t + w_t). \quad (1)$$

Taking logs,

$$\log(A_t) = s_t + \mu_t + w_t, \quad (2)$$

where s_t is a trend component, μ_t is the conditional mean of the detrended series and w_t is a mean zero error term. We assume that

$$w_t = v_t - u_t, \quad (3)$$

where u_t is an exponentially distributed (positive) random variable with expectation μ_u , giving information about the degree of asymmetry in the composite error, and v_t is a normally-distributed two-sided error, such that $v_t \sim N(\mu_u, \sigma_v^2)$. Note that the restriction $E(u_t) = E(v_t)$ ensures that $E(w_t) = 0$.³ The component s_t is removed from the data by applying the Hodrick-Prescott filter (Hodrick and Prescott 1997)⁴, and we consider the filtered series

$$y_t = \log(A_t) - s_t \quad (4)$$

as our dependent variable. The mean dynamics μ_t of the filtered series y_t are modeled by the AR(p) model

$$y_t = \alpha + \sum_{i=1}^p \beta_i y_{t-i} + w_t, \quad (5)$$

where the usual stationarity conditions are assumed to be satisfied and w_t is the composite error term specified above.⁵

³Other typical choices for the one-sided error term u_t are the half-normal, truncated normal, or gamma distribution (Kumbhakar and Lovell 2000, Murillo-Zamorano 2004). We initially considered the half-normal distribution for the one-sided error u_t . However, the exponential distribution fits the data substantially better; therefore, we present only this model.

⁴We used a smoothing parameter of 1600, the standard value for quarterly data. However, we checked the robustness of our results with respect to this choice.

⁵Initially, we also considered ARMA(p,q) models, but parsimonious AR models turned out to be sufficient to capture the dynamics in the data.

2.2 Estimation and evaluation

We estimate the parameters of the model, including the AR parameters and the parameters of the distributions of u_t and v_t , using maximum likelihood estimation. The lag length was chosen using a combination of information criteria and the Ljung-Box test for autocorrelation up to 24 lags. We allowed for the possibility of dropping insignificant intermediate autoregressive lags to reduce the number of parameters to be estimated. However, we checked the robustness of our results with respect to that choice.

A closed form expression for the density of the composite error $w_t = v_t - u_t$ exists and is given by (see Kumbhakar and Lovell 2000)

$$f(w_t; \mu_u, \sigma_v) = \frac{1}{\mu_u} \Phi \left(-\frac{w_t - \mu_u}{\sigma_v} - \frac{\sigma_v}{\mu_u} \right) \exp \left(\frac{w_t - \mu_u}{\mu_u} + \frac{\sigma_v^2}{2\mu_u^2} \right), \quad (6)$$

where Φ denotes the standard normal distribution function and μ_u is the parameter of the exponential distribution, which has mean equal to μ_u and variance equal to μ_u^2 . The mean of u_t is subtracted from w_t to account for the fact that, unlike in the classical stochastic frontier model, w_t is assumed to have mean zero in our situation. The log-likelihood function of the joint model, including the autoregressive dynamics, can be obtained in a straightforward way.

A primary goal of this paper is to test whether the model, allowing for one-sided error component in the AR innovations, outperforms the standard model with symmetric errors. The model with symmetric errors is nested in the model above when the parameter μ_u equals zero. Thus, one can test the null hypothesis that the benchmark model performs as well as our model by performing a likelihood ratio test. Lee (1993) showed that the likelihood ratio statistic asymptotically follows a mixture of a χ^2 distribution with one degree of freedom and a point mass of 1/2 at zero.

Estimation of this model is problematic when the sample skewness is positive. Aigner et al. (1977) demonstrated that theoretically in such situations the MLE of μ_u will converge to zero, and Lee (1993) showed that in this situation the information matrix is singular, which implies that maximum likelihood standard errors cannot be calculated. In practice, when the residuals have positive skewness, the MLE using (6) will either fail to converge or will converge to a local maximum. For cases in which the sample skewness of the residuals of the autoregressive model with Gaussian innovations is positive, we extend the model to explicitly allow for positive skewness. We allow $-u_t$ to follow an exponential distribution, such that u_t has a negative mean μ_u and standard deviation

$-\mu_u$. This leads to the composite density of w_t ,

$$\begin{aligned} f(w_t; \mu_u, \sigma_v) &= \frac{1}{\mu_u} \Phi \left(-\frac{w_t - \mu_u}{\sigma_v} - \frac{\sigma_v}{\mu_u} \right) \exp \left(\frac{w_t - \mu_u}{\mu_u} + \frac{\sigma_v^2}{2\mu_u^2} \right) I(\mu_u \geq 0) \\ &- \frac{1}{\mu_u} \Phi \left(\frac{w_t - \mu_u}{\sigma_v} + \frac{\sigma_v}{\mu_u} \right) \exp \left(\frac{w_t - \mu_u}{\mu_u} + \frac{\sigma_v^2}{2\mu_u^2} \right) I(\mu_u < 0), \end{aligned} \quad (7)$$

which for $\mu_u > 0$ corresponds to the one given in (6) with negative skewness, and for $\mu_u < 0$ is the mirrored version with positive skewness. $I(\cdot)$ denotes the indicator function. Densities are depicted in Figure 1. Substituting (7) for equation (6) avoids convergence problems of the estimation algorithm due to positive sample skewness. Moreover, it is possible to use a standard χ^2 distribution for a likelihood ratio test because the parameter μ_u is no longer on the boundary under the null hypothesis. Even if the estimate of μ_u is negative, it could be insignificant and support the evidence of a symmetric error distribution. Note that this model specification is a vehicle to test our null hypothesis of interest, rather than our model of interest, which remains (6). We use equation (7) when the estimation of equation (6) fails to converge.

The relative importance of u_t can be measured by computing the variance ratio (VR) of the two error components, expressed as

$$VR = \frac{\mu_u^2}{\sigma_v^2}, \quad (8)$$


which is estimated using the parameter estimates for the two error components. Note that the variance ratio measures the degree of asymmetry in our data. This is in contrast to the traditional stochastic frontier literature where the inefficiency of a firm is measured by *technical efficiency* (TE) defined by Battese and Coelli (1988) as $E[\exp(-u)|w]$. TE measures the distance from the efficient frontier or actual output divided by optimal efficient output. However, in our context TE cannot be interpreted, because the expectation of u is absorbed by the mean of v and consequently only the variance of u relative to the variance of v is relevant.

2.3 Alternative specifications and extensions

We also considered extensions and variations of the baseline stochastic frontier model. First, we considered allowing the one-sided error to follow a half-normal distribution. The goodness of fit was inferior to the one using the exponential distribution.

Second, we allowed u_t and v_t to be correlated, using the model proposed by Smith (2008). However, the model fit did not improve and simulations suggested that it is

Figure 1: Density function (7) with $\sigma_v = 1$. Solid curve: $\mu_u = 0$ (Gaussian), long dashed curve: $\mu_u = 1.5$ (negative skewness), short dashed curve: $\mu_u = -1.5$ (positive skewness)



extremely difficult to identify the correlation coefficient in small samples, which leads to strongly biased estimates of the variance ratio.⁶

For the modeling of the dynamics, we alternatively considered the parameter μ_u of the distribution of u_t to vary over time following a (transformation of a) first order autoregressive stochastic process. Although such a model can capture certain dynamics present in the data, modeling the dynamics in the data directly using AR models turned out to provide a superior fit to the data.

Finally, we considered a model where only the one-sided error component was persistent. By putting this specific structure on the one-sided error component it would be

⁶Simulation results are available upon request.

possible to extract the two error components from the data and not just their variance ratio. Again, this model showed promising results, but estimation is highly complex and therefore the dynamics have to be restricted to an AR(1) process. Furthermore, in such a model the asymmetric component artificially captures most of the overall variation since it picks up the autocorrelation in the data, biasing our inferences on asymmetry.

3 Empirical Implementation

3.1 The Data

The data set consists of seasonally-adjusted quarterly observations on constant-price GDP, investment, and employment for all OECD countries which had data dating back at least to the early 1980's. OECD does not have data on labor hours, a preferable measure for the labor input, or on capital stock. The sample ends in 2010:Q2 and has country-specific start dates (in parenthesis): Australia (1967:Q1), Canada (1961:Q1), France (1978:Q1), Germany (1970:Q1), Italy (1981:Q1), Japan (1980:Q1), Korea (1983:Q1), Norway (1978:Q1), Switzerland (1980:Q1), United Kingdom (1969:Q2) and United States (1955:Q1). The data does not constitute a panel because units of measurement for each country's output and investment differ since they are measured in country-specific units of 2005 GDP. Data from the Penn World Tables does adjust cross-country data to comparable units, using purchasing power parity measures of relative prices, but that data exists only at annual frequency. Since we are interested in business cycle properties, quarterly frequencies are essential. Therefore, we estimate eleven separate equations, decomposing the innovations of each country's Solow residual into a symmetric and an asymmetric component.

To construct Solow residuals, we first construct measures of the capital stock. We use the perpetual inventory method, letting the initial value of the capital stock (K_0) be the steady-state equilibrium value with the growth rate (g) equal to the average of growth over the first ten years of the sample, annual depreciation (δ) at 0.07, following Easterly and Levine (2001), and initial investment equal to its initial value (I_0).⁷ Subsequent values for capital are computed using the equation for the adjustment of the capital stock,

$$K_{t+1} = (1 - \delta) K_t + I_t \tag{9}$$

To compute Solow residuals, we use employment as the measure of labor input and set capital's share at 0.35, following Stock and Watson (1999).

⁷The initial capital stock becomes $K_0 = \frac{I_0}{g+\delta}$.

Figures 3-13 in the Appendix show the Solow residuals and the HP filtered series for all countries. Before estimating our models we standardize the (filtered) series to have variance equal to one.⁸ Empirical autocorrelations and partial autocorrelations (not reported) suggest that autoregressive models capture the dynamics in the data.

Data on oil production are available annually beginning in 1980 from the US Energy Information Administration (EIA), measured as thousand barrels per day. We use this data to compute our measure of average oil production per employed worker for each country. For oil prices we use the dollar spot price for West Texas Intermediate, and deflate this by the seasonally adjusted US CPI. The data are monthly, and we use end of period values for quarterly values.

3.2 Estimation results

Table 1 reports the results of the empirical analysis. It contains the estimated parameters of the model, the sample skewness of the residuals from the AR models using asymmetric errors, the log-likelihoods of both the model restricted to symmetry (LL AR) and our dynamic stochastic frontier model (LL DSF), the likelihood ratio statistic (LR stat.) for the null of symmetry along with its p-value, and the variance ratio (VR) implied by the estimated model. As mentioned in Section 2.2, when the residuals from the Gaussian AR model exhibit positive skewness we estimate the extended model (7) that allows for positive and negative skewness. This allows us to compute standard errors and perform the likelihood ratio test.

The likelihood ratio test shows that the symmetric model is significantly outperformed by the model allowing for asymmetries for all countries except Canada, Norway, and possibly France. France is a borderline case with a p-value of 0.108 and residuals with negative skewness. The UK has residuals with positive skewness, but the UK is a special case that we treat separately below. The estimated variance ratio for countries with evidence of negative asymmetries ranges from 0.28 for the US to 0.72 for Japan.

⁸The standardization does not affect the estimation of the quantities of interest such as the skewness or the variance ratio.

Table 1: Estimation results of the DSF model

	aus	can	fra	ger	ita	jap	kor	nor	swi	uk	uk (pre 94)	uk (post 94)	us
AR(1)	0.59***	0.84***	1.01***	0.69***	0.85***	0.74***	0.68***	0.51***	0.95***	0.85***	0.72***	1.05***	0.85***
AR(2)					-0.19***			0.25***					
AR(3)			-0.2***					-0.16**					
AR(4)	-0.19***	-0.12***		0.38***		-0.12**			-0.21***	-0.21***		-0.23***	-0.14***
AR(5)				-0.46***									0.10**
μ_u	0.47***	-0.26**	0.25***	0.44***	0.36***	0.43***	0.44***	-0.05	0.26***	-0.37***	-0.39***	0.27***	0.28***
σ_v	0.65***	0.56***	0.4***	0.56***	0.47***	0.50***	0.54***	0.78***	0.34***	0.45***	0.52***	0.26***	0.52***
skewness	-0.411	0.13	-0.283	-0.437	-0.475	-0.873	-0.952	0.030	-1.013	1.149	0.877	-1.364	-0.025
LL AR	-204.2	-182.9	-84.1	-174.4	-101.1	-121.1	-120.9	-148.9	-70.6	-150.6	-101.3	-29.9	-195.1
LL DSF	-201.2	-182.6	-83.4	-170.8	-98.9	-116.3	-114.1	-148.9	-65.2	-139.77	-96.4	-26.1	-194.2
LR stat.	6.074	0.5045	1.530	6.972	4.191	9.453	13.434	0.000	10.942	21.632	9.638	7.561	1.951
p-val.	0.007	0.4776	0.108	0.004	0.020	0.001	0.000	0.989	0.001	0.000	0.002	0.003	0.0812
VR	0.538	0.000	0.389	0.604	0.578	0.718	0.657	0.000	0.588	0.000	0.000	1.072	0.284

Note: Table 1 reports the estimation results of model (6) defined in Section 2.1 consisting of an AR model with an error term composed of a one-sided exponentially distributed term with parameter μ_u and a two-sided Gaussian term with parameter σ_v . The data are quarterly Solow residuals calculated as explained in Section 3.1 and detrended using the Hodrick-Prescott filter. Skewness refers to the sample skewness of the residuals from the model with asymmetric residuals, LL AR and LL DSF are the log-likelihood of the model restricted to symmetry and the general model respectively, LR stat. is the log likelihood ratio statistics for the null hypothesis of a symmetrically distributed error term, and VR refers to the ratio of the estimated variances of the one-sided and the two sided innovations. When sample skewness is positive, convergence of model (6) failed, and we present results from model (7). ***, ** and * refer to significance at the 1%, 5% and 10% confidence level, respectively.

It is interesting to consider which countries exhibit asymmetries. Countries which do not exhibit negative asymmetries are significant producers of oil. The UK has a large value of oil production per worker prior to 1994 and a very small value thereafter. Therefore, we separate the sample for the UK into two subsamples and reestimate our model for the two subsamples. The results in Table 1 show that the UK exhibits positive asymmetries prior to 1994, with some evidence for negative asymmetries after 1994.⁹ Thus, when we include the borderline case of France and pre-1994 UK, nine of the eleven countries have significant evidence of negative asymmetries.


In Figure 2, we compare variance ratios and mean oil production per worker from 1980 – 2009, setting variance ratios equal to zero whenever the asymmetric component is insignificant or positive. Countries with insignificant or positive asymmetries, Canada, Norway, and pre-1994 UK, have large values of oil production per worker. Countries with relatively large asymmetries, as measured by variance ratios, Germany, Italy, Japan, Korea, Switzerland, and post-1994 UK, have virtually no oil production per worker. Countries with intermediate values of oil per worker, the US and Australia, fall in between in their measure of variance ratios. France is a notable exception with very little oil and a variance ratio less than that of Australia.

Based on these results, we sought to determine whether oil price increases were responsible for the asymmetric negative components in Solow residuals for countries with significant asymmetries. We reestimated the model, conditioning on oil prices. The real oil price series are plotted in Figure 14 in the Appendix. Since the oil price series has a trend we again apply the HP filter to remove it, and we standardize the detrended data to have unit variance. Following Hamilton (2011), we allow linear and non-linear terms for oil prices and its lags. The non-linear terms are values for real oil price increases exceeding 1.5 (*Oil_pos*) and values for decreases less than 1.5 (*Oil_neg*).¹⁰ These results are contained in Table 2. Typically, either the linear term or the non-linear positive term are significant with the expected negative signs. In most cases, the first or second lag of *Oil_pos* leads to the best model fit. This extends Hamilton’s work for the US implying that oil price increases have asymmetrically large effects on Solow residuals for the US

⁹The post-1994 subsample ends with a very large positive residual. Excluding this outlier reduces the estimated sample skewness from 0.85 to -1.36. We present the variance ratio excluding this final observation, but realize that this is an upper bound on the true variance ratio. Excluding the final observation does not significantly affect the full sample results.

¹⁰We also allowed for different thresholds and we considered smooth transition regressions to allow for more general nonlinearities, but the results were robust with respect to the different specifications.

Figure 2: Oil production per worker vs. estimated variance ratio



and for eight other countries.¹¹ For Norway, a country with very large oil production per worker, the oil price is insignificant. We found that for seven of the nine countries with significant variance ratios, conditioning on oil prices reduces the asymmetries, as measured by the variance ratios.¹² The LR statistic shows that, in contrast to the results without oil, we cannot reject the symmetric model for Italy, the UK after 1994 and the US. For Japan we obtain significant positive asymmetry. Furthermore, all evidence for asymmetry disappears for France and for Germany with p-values increasing from 0.4% to slightly over 5%. Australia and Switzerland have asymmetries even after conditioning on oil prices. Korea retains asymmetries, but they are smaller. Inspection of the residuals for Korea reveals a large negative outlier in 1998:Q1, the precise timing of the Asian crisis. Conditioning on oil prices, together with a dummy for the Asian crisis, increases the skewness to 0.29, thereby eliminating all evidence of negative asymmetry.¹³

¹¹We note that Kilian and Vigfusson (2011) dispute Hamilton's finding.

¹²The variance ratios for Australia and Switzerland did not decrease.

¹³Conditioning only on the Asian crisis, the symmetric model is rejected with a p-value of 2.6%. In that case the skewness and variance ratio turn out to be equal to -0.15 and 0.351, respectively.

These results imply that changes in oil prices are responsible for at least some of the asymmetry in Solow residuals for most countries which have asymmetric negative components, equivalently those which produce little to moderate amounts of oil. Interestingly, even for countries which do not have a significant coefficient on *Oil_pos*, conditioning on oil prices reduces the variance ratios.

We checked the robustness of our results by using an AR(4) model with all intermediate lags included for all countries, except for the US and Italy where we went up to 8 lags. For the oil price we considered two alternative specifications. First we used four lags of *Oil_pos* and second we used four lags of each *Oil* and *Oil_pos*. The qualitative results of our analysis remained mostly unchanged. Furthermore, we also checked the robustness of our results with respect to the smoothing parameter of the HP filter. Again, the qualitative results were not changed.

Table 2: Estimation results of the DSF model including Oil prices

	aus	can	fra	ger	ita	jap	kor	swi	uk	uk (pre 94)	uk (post 94)	us
AR(1)	0.58***	0.73***	0.99***	0.66***	0.79***	0.76***	0.71***	0.95***	0.79***	0.69***	1.16***	0.83***
AR(2)					-0.15*							
AR(3)			-0.15***									
AR(4)	-0.18***	-0.07		0.42***		-0.12**		-0.18***	-0.21***		-0.25***	-0.15***
AR(5)				-0.43***								-0.10***
<i>Oil(t)</i>		-0.12***					-0.2**					
<i>Oil(t - 1)</i>			-0.16***									
<i>Oil_pos(t - 1)</i>												
<i>Oil_pos(t - 2)</i>	-0.16*			-0.34***	-0.26***	-0.48***			-0.14***	-0.15***	-0.26***	-0.22***
<i>Oil_pos(t - 3)</i>		-0.22***										
<i>Oil_neg(t - 2)</i>												
μ_u	0.48***	0.20	0.09	0.38***	0.29***	-0.34***	0.41***	0.25***	-0.36***	-0.38***	-0.16	-0.26***
σ_v	0.64***	0.56***	0.44***	0.58***	0.49***	0.51***	0.54***	0.33***	0.44***	0.51***	0.30***	0.51***
skewness	-0.424	-0.006	-0.063	-0.195	-0.219	0.582	-0.731	-1.207	1.065	0.923	0.160	0.073
LL AR	-202.9	-173.97	-77.86	-166.7	-95.6	-113.5	-115.0	-66.8	-144.8	-99.1	-21.8	-188.4
LL DSF	-199.8	-173.9	-77.86	-165.4	-95.1	-111.5	110.4	-60.2	-135.3	-93.9	21.7	-187.1
LR stat.	6.272	0.147	0.002	2.616	1.177	4.030	9.134	13.249	19.058	10.199	0.254	2.512
p-val.	0.006	0.351	0.482	0.053	0.139	0.045	0.001	0.000	0.000	0.001	0.614	0.113
VR	0.548	0.134	0.039	0.422	0.362	0.000	0.563	0.590	0.000	0.000	0.000	0.000

Note: Table 2 reports the estimation results of model (6) defined in Section 2.1 consisting of an AR model with an error term composed of a one-sided exponentially distributed term with parameter μ_u and a two-sided Gaussian term with parameter σ_v . Additionally, real oil prices are included as exogenous variables in the model. Oil pos. refers to oil prices exceeding the fixed threshold of 1.5 and Oil neg. refers to oil prices below -1.5. The data are quarterly Solow residuals calculated as explained in Section 3.1 and detrended using the Hodrick-Prescott filter. Skewness refers to the sample skewness of the residuals from the model with asymmetric residuals, LL AR and LL DSF are the log-likelihood of the model restricted to symmetry and the general model respectively, LR stat. is the log likelihood ratio statistics for the null hypothesis of a symmetrically distributed error term, and VR refers to the ratio of the estimated variances of the one-sided and the two sided innovations. When sample skewness is positive, convergence for model (6) failed, and we present results from model (7). ***, ** and * refer to significance at the 1%, 5% and 10% confidence level, respectively.

4 Conclusion

We apply a novel empirical technique to an important topic in business cycle analysis, business-cycle asymmetry. We investigate whether innovations in total factor productivity, measured as HP-filtered quarterly Solow residuals and considered an important determinant of business cycles, are symmetric across time for eleven OECD countries. Our model adapts stochastic frontier analysis to accommodate the strong serial correlation in macro data. The results have implications for whether business cycles are symmetric, with the economy responding in a linear way to normal iid shocks, or asymmetric, with recessions fundamentally different from booms. Log likelihood ratio tests imply that nine of eleven countries have significant negative asymmetries.

Second, we consider structural differences in economies with and without asymmetries and find that asymmetries tend to be stronger in countries with less oil production per worker. Non-linear conditioning of HP-filtered oil prices removes or reduces asymmetries for most countries which otherwise exhibit them, implying that much asymmetry is due to the response of the economy to oil prices. These results imply that positive and negative TFP shocks are caused by different fundamentals with different distributions.

These results raise several interesting issues for macroeconomic modeling of business cycles. The finding that Solow residuals have significant asymmetries implies that we are not going to understand business cycles in most countries if we continue to model them using iid productivity shocks. Our results argue for a fundamental modification of the standard DSGE paradigm in macroeconomics. Additionally, we need more work to understand the nature of the negative asymmetry. Our work extends and confirms Hamilton's (1983, 2011) work for the US, that one cause of the asymmetry is oil prices. We hypothesize that this could be the result of the response of bounded capacity utilization to demand shocks created by a terms of trade shift, but more research is needed to understand these fundamentals.

A Graphs

Figure 3: Solow residuals of Australia




Figure 4: Solow residuals of Canada




Figure 5: Solow residuals of France




Figure 6: Solow residuals of Germany




Figure 7: Solow residuals of Italy




Figure 8: Solow residuals of Japan






Figure 9: Solow residuals of Korea



(a) Raw data

(b) HP filtered data


Figure 10: Solow residuals of Norway



(a) Raw data

(b) HP filtered data

Figure 11: Solow residuals of Switzerland



(a) Raw data

(b) HP filtered data

Figure 12: Solow residuals of the United Kingdom




Figure 13: Solow residuals of the United States





Figure 14: Real oil price data



References

- Aigner, D. J., C. A. K. Lovel, and P. J. Schmidt (1977). Formulation and estimation of stochastic frontier production function models. *Journal of Econometrics* 6, 21–37.
- Basu, S. (1996). Procyclical productivity: Increasing returns or cyclical utilization? *Quarterly Journal of Economics* 111, 719–752.
- Battese, G. E. and T. J. Coelli (1988). Prediction of firm-level technical efficiencies with a generalized frontier production function and panel data. *Journal of Econometrics* 38, 387–399.
- Bils, M. and J. O. Cho (1994). Cyclical factor utilization. *Journal of Monetary Economics*, 319–354.
- Burnside, C. and M. Eichenbaum (1996). Factor hoarding and the propagation of business-cycle shocks. *American Economic Review* 86, 1154–1174.
- Burnside, C., M. Eichenbaum, and S. Rebelo (1993). Labor hoarding and the business cycle. *Journal of Political Economy* 101, 245–273.
- Evans, C. L. (1992). Productivity shocks and real business cycles. *Journal of Monetary Economics* 29, 191–208.
- Greenwood, J., Z. Hercowitz, and P. Krusell (2000). The role of investment-specific technological change in the business cycle. *European Economic Review* 44, 91–115.
- Hamilton, J. (1994). *Time Series Analysis*. Princeton University Press.
- Hamilton, J. D. (1983). Oil and the macroeconomy since world war ii. *Journal of Political Economy* 91, 228–248.
- Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica* 57, 357–384.
- Hamilton, J. D. (2011). Nonlinearities and the macroeconomic effects of oil prices. *Macroeconomic Dynamics*, forthcoming.
- Hodrick, R. and E. C. Prescott (1997). Postwar u.s. business cycles: An empirical investigation. *Journal of Money, Credit, and Banking* 29, 1–16.

- Kehoe, T. J. and K. J. Ruhl (2008). Are shocks to the terms of trade shocks to productivity? *Review of Economic Dynamics* 11, 804–819.
- Kilian, L. and R. J. Vigfusson (2011). Are the responses of the u.s. economy asymmetric in energy price increases and decreases? *Quantitative Economics*, forthcoming.
- Kumbhakar, S. C. and C. A. K. Lovell (2000). *Stochastic Frontier Analysis*. Cambridge University Press.
- Lee, L. F. (1993). Asymptotic distribution of the maximum likelihood estimator for a stochastic frontier function model with a singular information matrix. *Econometric Theory* 9, 413–430.
- Meeusen, W. and J. van den Broek (1977). Efficiency estimation from cobb-douglas production functions with composed errors. *International Economic Review* 18, 435–444.
- Murillo-Zamorano, L. R. (2004). Economic efficiency and frontier techniques. *Journal of Economic Surveys* 18, 33–77.
- Smith, M. D. (2008). Stochastic frontier models with dependent error components. *Econometrics Journal* 11, 172–192.