Multiple filtering devices for the estimation of cyclical DSGE models

FABIO CANOVA ICREA-UPF, CREI, CREMeD, and CEPR

FILIPPO FERRONI Banque de France

We propose a method to estimate time invariant cyclical dynamic stochastic general equilibrium models using the information provided by a variety of filters. We treat data filtered with alternative procedures as contaminated proxies of the relevant model-based quantities and estimate structural and nonstructural parameters jointly using a signal extraction approach. We employ simulated data to illustrate the properties of the procedure and compare our conclusions with those obtained when just one filter is used. We revisit the role of money in the transmission of monetary business cycles.

KEYWORDS. DSGE models, filters, structural estimation, business cycles. JEL classification. C32, E32.

1. Introduction

Dynamic stochastic general equilibrium (DSGE) models have become the paradigm for business cycle and policy analyses in academic and policy circles. Relative to earlier structures, current models are of larger scale and feature numerous real and nominal frictions that help to closely replicate the dynamic responses that structural vector autoregressions (VARs) produce. A few years ago it was standard to informally calibrate these models, but today, increased computing power and recent developments in systemwide estimation methods allow researchers to routinely employ full information techniques in structural estimation exercises.

Despite the increased popularity, structural estimation faces important conceptual and numerical problems. For example, as emphasized in Canova (2009), full informa-

Fabio Canova: fabio.canova@upf.edu

Filippo Ferroni: filippo.ferroni@banque-france.fr

We thank two anonymous referees and an editor of the journal for detailed comments, and Don Harding, Adrian Pagan, Anton Braun, Vasco Carvalho, Kris Nimark, the participants of the Central Bank workshop "Macroeconomic Modeling 2008," Cartagena, Colombia, of the CEMLA meeting of researchers, Lima, Peru, and of seminars at CREI, the Czech National Bank, BIS and the Australian Reserve Bank, and Bank of Italy for suggestions. The financial support of the Spanish Ministry of Education, through the Grant SEJ-2006-02235, the Spanish Ministry of Science and Technology, through the Grant ECO2009-08556, and of the Barcelona Graduate School of Economic is gratefully acknowledged. The views presented in this paper do not reflect those of the Banque de France.

 $Copyright © 2011 \ Fabio \ Canova \ and \ Filippo \ Ferroni. \ Licensed \ under the \ Creative \ Commons \ Attribution-NonCommercial \ License \ 3.0. \ Available \ at \ http://www.qeconomics.org.$

DOI: 10.3982/QE36

tion classical estimation makes sense only if the model is the data generating process (DGP) of the observables, up to a set of serially uncorrelated measurement errors. Since such an assumption is hard to entrain unless the model is augmented with ad hoc dynamics, Fukac and Pagan (2010) suggested to complement standard inference with a more robust limited information analysis. It is also well known that there are abundant population identification problems (see Canova and Sala (2009), Del Negro and Schorfheide (2008)), that numerical difficulties are widespread, and that errors-invariables are present (the variables in the model do not often have a direct counterpart in the data). Finally, the vast majority of the models used in the literature are time invariant and intended to explain only the cyclical portion of the observable fluctuations, while the actual data contain many types of fluctuations, all of which may be subject to breaks and other forms of slowly moving variations.

To fit stationary cyclical DSGE models to the data, applied investigators typically select a subsample where time invariance is more likely to hold, filter the raw data with an arbitrary statistical device, and treat the filtered data as the relevant measure of stationary cyclical fluctuations (see, e.g., Smets and Wouters (2003), Ireland (2004)). Alternatively, one arbitrarily builds a noncyclical component into the model (e.g., via a deterministic labor augmenting technology progress or unit roots in total factor productivity and/or the price of investment) and filters the raw data using a model-driven transformation (see, e.g., Fernandez-Villaverde and Rubio-Ramirez (2007) or Justiniano, Primiceri, and Tambalotti (2011)) or an arbitrary statistical device (see Smets and Wouters (2007)).

Both approaches are, in general, problematic. While the profession shares the idea that a cyclical model should explain fluctuations with an average periodicity of 8-32 quarters, there is little agreement on how to obtain these fluctuations from the data and only a partial understanding of the consequences that statistical filtering induces. For example, it is common to use linearly detrended or first differenced data as input in the estimation process, but such transformations do not isolate fluctuations with the required periodicity (see, e.g., Canova (1998)). A band pass (BP) filter, which can potentially extract the fluctuations of interest with an infinite amount of data, is typically discarded in the estimation literature because its two-sided nature alters the timing of the data information; a similar argument is also made for the Hodrick and Prescott (HP) filter. Moreover, while real variables typically show long run drifts, nominal variables just display low frequency fluctuations. Hence, should we filter all the data or only real variables? Investigators have taken both positions, but it is not obvious which approach is preferable. Finally, since researchers filter each series separately, theoretically relevant constraints may not be satisfied with filtered data (for example, does a resource constraint hold with filtered data?).

Model-driven filtering also fails to extract cycles with the required periodicity. For example, when total factor productivity (TFP) is trending, real variables share similar trends and appropriate linear combinations should be free of noncyclical dynamics. However, as shown in Canova (2008), real and nominal "great ratios" display significant upward drifts and the portion of the variance of the transformed variables located outside the cyclical frequencies is generally large. Most problematic of all, model-based filtering requires knowledge of the number, the nature, and the time series features of the

shocks driving the noncyclical component. Given our general ignorance on the subject, important specification errors may plague structural estimates.

Since solving this complex mismatch problem is difficult, this paper focuses on how to improve structural estimation of the parameters of a cyclical DSGE model when a statistical filtering approach is used to match the data to the model counterparts. We make three contributions to the existing literature. First, we show that a typical log-linearized DSGE model produces cyclical fluctuations which are not necessarily located at the socalled business cycle frequencies. Thus, standard filtering approaches induce measurement errors in the estimated cyclical components. Since these errors have important low frequency components, the true income and substitution effects are mismeasured, leading to distortions in the estimates of important structural parameters. Second, we show how to design a statistical filter which captures the cyclical component of a DSGE model. This filter is model specific and the computational complexities involved make its practical implementation unfeasible on current computers. Third, we propose a method to estimate the structural parameters of a time invariant cyclical DSGE model which may potentially eliminate the biases that statistical filters produce. The approach borrows ideas from the recent data-rich environment literature (see Boivin and Giannoni (2005)). We set up a signal extraction framework where the cyclical DSGE is the unobservable factor; vectors of filtered data are contaminated observable proxies, and DSGE and nonstructural parameters are jointly estimated.

Our approach is advantageous in at least two respects. Since we do not have to arbitrarily choose one filtering method prior to the estimation or select which shock drives the noncyclical component, we avoid important specification errors. Moreover, our method can be used with cyclical data, which are obtained with one-sided and twosided filters, of both univariate and multivariate nature, as long as the list of filters is sufficiently rich. For the approach to work properly, the list of filters should be carefully chosen and suggestions on how to do this in practice are provided.

We investigate the properties of our approach using experimental data of the typical length employed in macroeconomics and demonstrate that the biases obtained when just one filter is used are reduced with our approach. We also show that the unconditional one-step-ahead mean square error (MSE) produced by our approach is smaller than the MSE obtained with standard procedures and that conditional forecasts are better behaved.

To show that the biases are also economically relevant, we revisit the role of money in amplifying cyclical fluctuations. The recent literature has neglected the stock of money when studying monetary business cycles, and Ireland (2004) demonstrated that such an approach is, by and large, appropriate using U.S. data, standard filtering techniques, and a maximum likelihood estimator. We show that when multiple filtered data are jointly used in the estimation, money balances matter for the transmission of cyclical fluctuations to output and inflation, and the propagation of primitive shocks differs from the one obtained when only one data transformation is used.

We want to be clear why we insist on working with time invariant cyclical models, rather than considering structures where cyclical and noncyclical fluctuations are jointly accounted for. On one hand, constructing reasonable models with these features is hard: theory is largely silent on how cyclical shocks can be propagated at longer frequencies (exceptions are Comin and Gertler (2006) or Michelacci and Lopez-Salido (2007)) or how long run disturbances can produce important cyclical implications. Moreover, it is convenient for both policy and interpretation purposes to assume that the mechanisms driving cyclical and noncyclical fluctuations are distinct and orthogonal. Finally, breaks make the data largely uninformative about the features of noncyclical fluctuations.

The rest of the paper is organized as follows. The next section shows the problems one encounters using a single filter to estimate the parameters of DSGE models. Section 3 derives the features of an optimal filter. Section 4 presents our approach. Section 5 examines the role of money in transmitting monetary business cycles. Section 6 concludes. The Appendixes are available in the Supplemental Material (Canova and Ferroni (2011)).

2. STATISTICAL FILTERS AND STRUCTURAL PARAMETER ESTIMATES

To show that statistical filtering induces important measurement errors in the estimated cyclical components and to investigate how these errors affect structural estimates, we simulate data from a textbook new-Keynesian model (see, e.g., Gali (2008)), where agents face a labor–leisure choice, production is carried out with labor, firms face an exogenous probability of price adjustments, and monetary policy is represented with a conventional Taylor rule. The equilibrium conditions are

$$0 = \chi_t (C_t - hC_{t-1})^{-\sigma_c} - \mathcal{L}_t, \tag{1}$$

$$0 = N_t^{\sigma_n} - \mathcal{L}_t \frac{W_t}{P_t},\tag{2}$$

$$1 = E_t \left[\beta \frac{\mathcal{L}_{t+1}}{\mathcal{L}_t} \frac{R_t}{\Pi_{t+1}} \right], \tag{3}$$

$$0 = E_{t} \left(\sum_{k=0}^{\infty} \frac{P_{t+k} \mathcal{L}_{t+k}}{P_{t} \mathcal{L}_{t}} (\beta \zeta_{p})^{k} \right) \times \left[1 - \epsilon_{t+k} + MC_{t+k}^{r} \epsilon_{t+k} \left(\frac{\widetilde{P}_{t}}{P_{t+k}} \right)^{(\alpha - 1 - \alpha \epsilon_{t+k})/(1 - \alpha)} \right] Y_{t+k}(j) \right),$$

$$(4)$$

$$1 = \left(\zeta_p \left(\frac{P_{t-1}}{P_t}\right)^{1-\epsilon_t} + (1-\zeta_p) \left(\frac{\widetilde{P}_t}{P_t}\right)^{1-\epsilon_t}\right)^{1/(1-\epsilon_t)},\tag{5}$$

$$Y_t = C_t, (6)$$

$$N_t = \left(\frac{Y_t}{Z_t}\right)^{1/(1-\alpha)} \int_0^1 \left(\frac{P_t(j)}{P_t}\right)^{-\epsilon_t/(1-\alpha)} dj,\tag{7}$$

$$MC_t^r = \frac{W_t}{P_t} \left(\frac{1}{Z_t}\right)^{1/(1-\alpha)} Y_t^{\alpha/(1-\alpha)} \int_0^1 \left(\frac{P_t(j)}{P_t}\right)^{-\epsilon_t \alpha/(1-\alpha)} dj, \tag{8}$$

$$R_t = R_{t-1}^{\rho_r} \pi_t^{(1-\rho_r)\rho_\pi} Y_t^{(1-\rho_r)\rho_y} v_t, \tag{9}$$

where h is the consumption habit coefficient, σ_c is the risk aversion coefficient, $1/\sigma_n$ is the Frisch elasticity, β is the discount factor, $1 - \alpha$ is the share of labor in production, $1-\zeta_p$ is the probability of changing prices, and ρ_{π} , ρ_{ν} , and ρ_r are the parameters of the monetary policy rule; \mathcal{L}_t is the Lagrangian on the consumer budget constraint, Y_t is aggregate output, $Y_t(j)$ is output of good j, N_t is aggregate hours, W_t is the nominal wage, R_t is the nominal interest rate, π_t is the inflation rate, P_t is the price level, $P_t(j)$ is the price of good j, MC_t^r is aggregate real marginal costs, and \widetilde{P}_t is the optimal price; χ_t is a preference shock, Z_t is a technology shock, ϵ_t is a markup shock, and v_t is a monetary policy shock. The first equation equates the marginal utility of consumption to the Lagrangian; the second equation relates the intertemporal rate of substitution between leisure and consumption to the real wage, and the third equation is a pricing relationship for one period real bonds. The next equation is a Phillips curve. Equation (5) describes the behavior of the aggregate price level. Equations (6), (7), and (8) define the resource constraints, aggregate hours, and real marginal costs. The last equation is the policy rule of the central bank. A full description of the model and the log-linearized conditions are given in Appendix A.

For the sake of illustration, we consider two situations. In the first one, $\ln \chi_t =$ $\rho_{\chi} \ln \chi_{t-1} + e_t$, where $e_t \sim N(0, \sigma_{\chi}^2)$; $\ln \epsilon_t = \epsilon + \frac{1-\epsilon}{\epsilon} \mu_t$, where $\mu_t \sim N(0, \sigma_{\mu}^2)$, $\ln v_t \sim 1$ $N(0, \sigma_v^2)$, and $Z_t = Z_{t,c}Z_{t,T}$, where $\ln Z_{t,T} = \gamma t + e_{t,T}$ with $e_{t,T} \sim N(0, \sigma_{Z,T}^2)$ and $\ln Z_{t,c} = \rho_z \ln Z_{t-1,c} + e_{t,c}$ with $e_{t,c} \sim N(0, \sigma_{Z,c}^2)$ (DGP1). In the second case $\chi_t = \chi_{t,c} \chi_{t,T}$, where $\ln \chi_{t,c} = \rho_{\chi} \ln \chi_{t-1,c} + e_{t,c}$ with $e_{t,c} \sim N(0, \sigma_{\chi,c}^2)$; $\ln \chi_{t,T} = \ln \chi_{t-1,T} + e_{t,T}$ with $e_{t,T} \sim N(0, \sigma_{\chi,T}^2)$; $\ln \epsilon_t = \epsilon + \frac{1-\epsilon}{\epsilon} \mu_t$, where $\mu_t \sim N(0, \sigma_\mu^2)$, $\ln v_t \sim N(0, \sigma_v^2)$, and $\ln Z_t = \frac{1-\epsilon}{\epsilon} \mu_t$ $\rho_z \ln Z_{t-1} + e_t$, where $e_t \sim N(0, \sigma_Z^2)$ (DGP2). Thus, in both specifications, there are four shocks driving cyclical (stationary) fluctuations and one shock driving noncyclical (nonstationary) fluctuations. However, in DGP1, noncyclical fluctuations are driven by a technology shock which is stochastic around a linear trend; in DGP2, they are driven by a preference shock that displays a unit root. For both DGPs, we set $\beta = 0.99$, $\sigma_c = 1.00, h = 0.70, \sigma_n = 0.70, \epsilon = 7.0, \rho_r = 0.2, \rho_{\pi} = 1.30, \rho_y = 0.05, \zeta_p = 0.8, \rho_{\chi} = 0.5,$ $\rho_z = 0.8$, $\sigma_v = 0.0012$, and $\sigma_\mu = 0.2064$. In DGP1, we select $\alpha = 0.4$, $\sigma_\chi = 0.0112$, $\gamma = 0.002$, $\sigma_{Z,T} = 0.003$, and $\sigma_{Z,c} = 0.0051$; in DGP2, $\alpha = 0.0$; $\sigma_{Z} = 0.0051$, $\sigma_{\chi,c} = 0.0112$, and $\sigma_{\chi,T} = 0.0012$. None of the points we make, however, depends on the choice of these parameters.

Table 1 presents a few moments of filtered output and filtered inflation when linear (LT), Hodrick and Prescott (HP), band pass (BP), and first order difference (FOD) filtering are used together with the moments of their true cyclical component, when T = 1000 this sample size effectively reduces small sample biases to zero. Clearly, regardless of the DGP, the variability, the serial, and the cross-correlation properties of the cyclical component of output and inflation are distorted. Also, although output displays a linear trend under DGP1 and a unit root under DGP2, LT filtering in DGP1 and FOD filtering in DGP2 are as biased as other arbitrary filtering approaches. Thus, misspecification of the noncyclical component cannot be the reason for these distortions. Finally, although DGP2 features a unit root, the raw inflation series is persistent but stationary. Hence, it will matter for structural estimation whether the model is fitted to filtered or unfiltered inflation.

DGP1 DGP2 Variable Filter St. Dev. AR(1) $corr(y, \pi)$ St. Dev. AR(1) $corr(y, \pi)$ Output LT 0.0486 0.864 0.925 0.0123 0.911 -0.196HP 0.0366 0.876 0.834 0.0065 0.691 -0.288BP 0.0377 0.908 0.939 0.0060 0.859 -0.553FOD 0.0188 0.608 0.513 0.0052 0.100 -0.0290.728 -0.324True 0.0295 0.914 0.0082 0.811 LT Inflation 0.0043 0.703 0.0100 -0.005HP 0.0037 0.602 0.0095 -0.083BP 0.0034 0.873 0.0050 -0.810**FOD** 0.0033 -0.1380.0138 -0.4950.590 True 0.0022 0.0098 0.005

Table 1. Moments of filtered and true cyclical components: simulated data.^a

To show how filtering errors affect parameter estimation, we take the experimental data for output, real wages, interest rates, and inflation constructed with DGP2 and estimate the structural parameters by prefiltering the raw data with LT, HP, BP, and FOD filters. Estimation is conducted with Bayesian methods: we choose relatively loose priors for all the parameters and, to give the routine the best chance, we start estimation at the true parameter values. Posterior estimates are obtained with a random walk Metropolis algorithm, where the jumping variable has a t-distribution with 5 degrees of freedom and the variance is tuned to have an acceptance rate of about 30 percent for each filtering approach. Half a million draws were made in each case; convergence was checked with a standard CUMSUM statistic and achieved after less than 250,000 iterations. We keep 1 out of 100 of the last 100,000 draws to compute posterior statistics. Results obtained with a flat prior are available on request from the authors. Table 2 reports the median and the standard deviation of the posterior of each structural parameter when all observables are independently filtered prior to estimation. Appendix B contains estimates for other relevant cases and other DGPs. There are important estimation biases in all cases and the magnitude of the bias can exceed 100 percent for some parameters. Interestingly, the parameters that regulate the relative magnitude of income and substitution effects (the Frisch elasticity σ_n^{-1} , the habit parameter h, the policy parameter ρ_{π} , and the persistence of the shocks) are considerably distorted. Estimates of the structural parameters appear to be relatively similar across three of the columns, but this outcome depends on the features of the DGP, in particular, on whether the noncyclical component is driven by technology or preference disturbances, on the relative variability of the noncyclical shocks, and on whether all observables or only a portion of them are filtered prior to estimation (see Appendix B).

While we have chosen to perform estimation using 150 data points to mimic a realistic estimation situation, larger samples will not change the conclusions. Thus, distortions obtain because of "population" rather than "small sample" errors. Similarly, allowing for measurement errors in the estimation will not change the features of Table 2:

^aAll variables are filtered prior to estimation. The sample size is T = 150.

TABLE 2. Parameter estimates obtained using different filters. The DGP features a preference shock with two components: a stationary AR(1) and a unit root.^a

				Filter			
				LT	HP	FOD	BP
	True	Prior [M	ean and s.d.]	Median (s.e.)	Median (s.e.)	Median (s.e.)	Median (s.e.)
σ_c	1.00	$\Gamma(0.1, 0.1)$	[1.00, 10.0]	3.77 (0.25)	4.38 (0.36)	2.21 (0.16)	5.23 (0.24)
σ_n	0.70	$\Gamma(0.5, 0.5)$	[1.00, 4.0]	0.28 (0.05)	0.13 (0.02)	0.04 (0.00)	0.06 (0.01)
h	0.70	B(10, 3)	[0.76, 0.11]	0.58 (0.03)	0.61 (0.06)	0.69 (0.03)	0.85 (0.05)
ϵ	7.00	N(6, 0.5)	[6.00, 0.50]	3.95 (0.13)	3.95 (0.13)	4.05 (0.13)	3.96 (0.13)
ρ_r	0.20	B(10, 6)	[0.71, 0.09]	0.30 (0.01)	0.27 (0.01)	0.39 (0.01)	0.59 (0.02)
ρ_{π}	1.30	N(1.5, 0.2)	[1.50, 0.20]	1.71 (0.06)	1.60 (0.05)	1.79 (0.06)	1.50 (0.05)
ρ_y	0.05	N(0.4, 0.2)	[0.40, 0.20]	-0.03(0.01)	-0.12(0.03)	0.01 (0.00)	-0.04(0.01)
ζ_p	0.80	B(6, 6)	[0.50, 0.14]	0.83 (0.03)	0.82 (0.03)	0.80 (0.03)	0.93 (0.03)
ρ_{χ}	0.50	B(10, 6)	[0.71, 0.09]	0.61 (0.03)	0.33 (0.02)	0.62 (0.05)	0.61 (0.04)
ρ_z	0.80	B(10, 6)	[0.71, 0.09]	0.72 (0.04)	0.54 (0.04)	0.24 (0.03)	0.70 (0.03)
$\sigma_{\chi,c}$	1.11	$\Gamma^{-1}(10, 20)$	[0.0056, 0.0020]	0.14 (0.02)	0.18 (0.16)	0.21 (0.05)	0.23 (0.43)
σ_z	0.51	$\Gamma^{-1}(10, 20)$	[0.0056, 0.0020]	0.15 (0.03)	0.27 (0.04)	3.87 (0.42)	1.72 (0.22)
σ_v	0.12	$\Gamma^{-1}(10, 20)$	[0.0056, 0.0020]	0.03 (0.00)	0.03 (0.00)	0.03 (0.00)	0.03 (0.00)
σ_{μ}	20.64	$\Gamma^{-1}(10,20)$	[0.0056, 0.0020]	7.31 (0.35)	4.90 (0.39)	4.96 (0.19)	5.77 (0.23)

^aAll variables are filtered prior to estimation. The sample size is T = 150. Γ stands for the gamma distribution, B for the beta distribution, and N for the normal distribution. In square brackets are the mean and the standard deviation of the prior.

the variability of the structural shocks is altered, but the magnitude and the direction of the biases in the estimates of important structural parameters are unchanged (for both exercises, see Appendix B).

To understand why distortions occur, it is useful to plot the spectral density of the cyclical component of output and inflation (obtained by setting $\gamma = \sigma_{z,T} = 0$ for DGP1 or $\sigma_{x,T} = 0$ for DGP2 in the simulations) together with the spectral density of the four filtered data when T = 1000. If one filtering transformation recovers the true cyclical component, the difference between the two spectra will be zero at all frequencies. Imperfect isolation in certain frequency bands will be evident when the two spectra differ considerably in those bands. To facilitate the discussion, we divide the spectrum into low, business cycle, and high frequencies, and, in Figure 1, we separate the frequencies that correspond to cycles of 8-32 quarters from the others with two vertical bars. Two observations are immediate. First, the cyclical component produced by a DSGE model does not have power only at the so-called business cycle frequencies; in fact, its spectrum resembles that of an AR(1) process. For the standard shock processes we have used, about half of the variability of the series is located at frequencies that correspond to cycles larger than 32 quarters. Thus, the idea that a statistical filter defines what is relevant for the analysis is incompatible with the assumption that a class of stationary DSGE models has generated the data. Moreover, focusing on business cycle frequencies is restrictive and may bias the interpretation of the economic phenomena. Second, even with 1000 data points, all filters imperfectly capture the spectrum of the true cyclical component of output and inflation. More importantly, regardless of the DGP, the filtering error is not

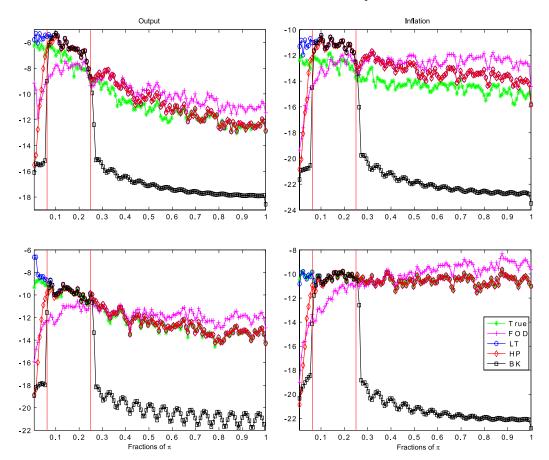


FIGURE 1. Log spectrum: true and estimated cyclical components. Top panel DGP1; bottom panel DGP2.

located only in the high frequencies and its frequency distribution is somewhat filter dependent. For example, LT filtered data have a stronger low frequency component and the other three filtered data have a weaker low frequency component than the actual cyclical data. At business cycle frequencies, the cyclical component extracted with HP, BP, and LT filters overestimates the true cyclical component of both variables, while FOD filtered data grossly underestimate the variability of the true cyclical component.

Why are errors present? The statistical filters we consider look like high pass or band pass filters. Thus, they appropriately extract the cyclical component of the data if and only if the noncyclical component of the model is solely located at those frequencies that are suppressed by the filters and the cyclical component is entirely located at the frequencies where the gain function of the filter is unity. Given that the cyclical component generated by a (log-linear) DSGE model will typically have power at all frequencies of the spectrum, filtering errors are created. In particular, since all the filters but LT attribute the power in low frequencies to the noncyclical component, important downward distortions are created at these frequencies. For the LT filter instead, upward distortions are produced because the stochastic elements of the noncyclical component

are disregarded. Mismeasurement of the low frequencies portion of the cyclical fluctuations is particularly troublesome because the estimated income and substitution effects are different from the true income and substitution effects, and this affects structural parameter estimates.

Knowing the DGP of the data is not a precondition for the above argument to hold. The (log-linear) solution of a stationary DSGE model has either an autoregressive (AR) or an autoregressive moving average (ARMA) representation (depending on whether all or a subset of the endogenous variables is considered), regardless of its exact structure. If the shocks are persistent, as it is usually assumed, it will always be the case that the data simulated by the stationary solution will have power in the low frequencies of the spectrum. Furthermore, the proportion of the variability in those frequencies is an increasing function of the persistence of the shocks. Our point is also completely independent of the assumed process that drives the noncyclical component and of its exact location (compare, e.g., the plots obtained with DGP1 and DGP2).

It is common among practitioners to believe that different filters are simply different ways to capture what generates the noncyclical component of the data. This perception is, in general, incorrect. The choice of filter also has implications for what we believe the cyclical component is. Incorrect filtering distorts both components and, as the discussion following Table 1 demonstrates, misspecification of the cyclical component may have more severe consequences than misspecification of the noncyclical component.

3. An ideal filter for DSGE models

To eliminate the distortions induced by imperfect filtering, one should design a filter that exploits the information that the cyclical components of a DSGE have the features of an AR (ARMA) process.

For this purpose, suppose a time series y_t has two components: c_t , which carries relevant information about the parameters of the model, and T_t , a nuisance component, and suppose for simplicity that T_t and c_t are uncorrelated (which, in our context, means that they are driven by independent shocks). Suppose we have available a time invariant linear filter g(L) and let $y_t^f = g(L)y_t$ be the filtered series. Under what conditions would $y_t^f = c_t$? For this to happen, we need $g(L)T_t = 0$ and $g(L)c_t = c_t$. In frequency domain, these two conditions imply that $g(\omega)(S_T(\omega) + S_c(\omega)) = S_c(\omega)$, where $S_i(\omega)$ is the spectral density of i=T, c at frequency ω and $g(\omega)$ is the square modulus of the transfer function of g(L). Then, if $0 \le g^*(\omega) = \frac{S_c(\omega)}{S_T(\omega) + S_c(\omega)} \le 1$, $y_t^f = c_t$. Thus, $g^*(\omega)$ needs to be large (small) at the frequencies where $S_c(\omega)$ is large (small).

In time series analysis, it is typical to assume that c_t has power only at certain frequencies, say, $S_c(\omega) \neq 0 \ \forall \omega \in (\omega_1, \omega_2)$ and that T_t has power at other frequencies, so that $S_T(\omega) = 0 \ \forall \omega \in (\omega_1, \omega_2)$. In this case, a band pass filter $g^*(\omega) = 1 \ \forall \omega \in (\omega_1, \omega_2)$ and $g^*(\omega) = 0$ otherwise, will make $y_t^f = c_t$. However, if c_t also has power for $\omega \in (0, \omega_1)$, $g^*(\omega) \neq 0 \ \forall \omega \in (0, \omega_1)$, a band pass filter fails to recover the true cyclical component.

As discussed, in log-linearized DSGE models, c_t has roughly the structure of a persistent AR (or a persistent ARMA) process, meaning that $S_c(\omega) \neq 0 \ \forall \omega$ and $\frac{\partial S_c(\omega)}{\partial \omega} < 0 \ \forall \omega$ (or $\frac{\partial S_c(\omega)}{\partial \omega} \leq 0$ for $\omega < \omega_1$ and $\frac{\partial S_c(\omega)}{\partial \omega} \geq 0$ for $\omega \geq \omega_1$). Hence, regardless of the exact structure of T_t , band pass or high pass filters (such as the HP or FOD filters) will induce measurement error at some or all frequencies. High frequency measurement error affects the standard errors of the estimates, but, in general, will not change the properties of point estimates (compare, for example, BP and HP estimates in Table 2). On the other hand, low frequency measurement errors are problematic.

Since $S_c(\omega) \neq 0 \ \forall \omega$ and $\frac{\partial S_c(\omega)}{\partial \omega} < 0 \ \forall \omega$, the ideal filter for a DSGE model must be such that $\log g^*(\omega)$ is increasing in ω , never vanishes over $(0, \pi)$, and approaches 1 only for $\omega = \pi$. If it is unique, it can be calculated with an iterative approach, which we summarize next:

Algorithm 3.1.

- Step 1. Choose a θ^0 vector of structural parameters and compute $c_t(\theta^0)$ using the model.
- Step 2. Given an observable y_t , obtain $g(\omega)(\theta^0) = \frac{S_c(\omega, \theta^0)}{S_y(\omega)}$ and compute $y_t^f(\theta^0) = g(L)(\theta^0)y_t$ using the resulting g(L) filter.
- Step 3. Use $y_t^f(\theta^0)$ to estimate the parameters of the model. Call the estimated vector θ^1 .
- Step 4. Iterate on Steps 1–3 until $||g(L)(\theta^i)y_t c_t(\theta^i)|| < \varepsilon$ or $||\theta^i \theta^{i-1}|| < \varepsilon$, or both, $i = 2, \ldots$

The metric in Step 4 is chosen by the investigator and may be frequency specific. Under the assumption that the data have been generated by a Markov process which is irreducible, aperiodic, and Harris recurrent, and that the metric used is the total variation norm over frequencies, adaptation of the results of Tierney (1994) will insure that convergence occurs as the number of iterations becomes large.

A few points about the algorithm are worth emphasizing. First, the optimal $g^*(\omega)$ does not necessarily generate a one-sided g(L) or weights g_j that decay fast to zero. Therefore, practical issues concerning alteration of the timing of the information and truncation need to be address. Second, the iterative procedure is time consuming since the model needs to be estimated numerous times before the fixed point is found. Given the computational costs of estimating the parameters of DSGE models by full information methods, this iterative approach is unfeasible on current computers. Finally, the ideal filter is model specific—it depends on what shocks drive c_t , their time series structure, and the features of the internal propagation mechanism of shocks—and it is subject to standard specification errors if the cyclical component is misspecified.

Given these difficulties, rather than trying to construct the ideal filter for a particular model, we prefer to take another route to improve the quality of the estimates of the structural parameters of a cyclical DSGE, which is not model specific and is computationally feasible. Our idea is to use the information contained in the cyclical data generated by a number of filters in the estimation. In practice, we treat cyclical data extracted with various filtering methods as contaminated estimates of the unobservable model-based cyclical component and use the information provided by a carefully selected list

of filters jointly in the estimation of the structural parameters. If the measurement error is close to being idiosyncratic across filtering methods in the low frequencies, our signal extraction approach will average it out. Thus, we obtain more precise estimates of the cyclical features of the economy and, hopefully, better estimates of the structural parameters are obtained.

How do we obtained improved estimates of the structural parameters? Let $g_i(L)$, $i=1,2,\ldots,q$, be different filters and let y_t^i be the resulting filtered data. Let $S_i(\omega)$, $i=1,2,\ldots,q$, be the spectral density of the filtered data y_t^i and assume that $S_i(\omega)=$ $S_c(\omega) + S_{u_i}(\omega)$. Then $\sum_i s_i(\omega) S_{u_i}(\omega) = 0$ for some set of weights $s_i(\omega)$ as long as $S_{u_i}(\omega)$ are sufficiently idiosyncratic across i. For weights $s_i(\omega)$ which are independent of ω , it may not be possible to set $\sum_{i} s_{i} S_{u_{i}}(\omega) = 0$ at all ω and one can choose them to reduce, for example, low frequency measurement error. Note that the use of cross-sectional information to identify the model-based cyclical component allows the u_i to be serially correlated. Given that q is finite here, we will not allow cross-filter correlation in the u_i and this requires a careful selection of filtering methods to be used in the estimation.

4. An alternative framework

Let the log-linearized solution of a cyclical DSGE model be

$$x_t = \Phi(\theta)x_{t-1} + \Psi(\theta)e_t, \quad e_t \sim (0, \Sigma(\theta)), \tag{10}$$

where Ψ and Φ are time invariant functions of the vector of structural parameters $\theta =$ $(\theta_1,\ldots,\theta_k)$, x_t are the endogenous variables, and e_t are the structural innovations. We let $x_t^m = Sx_t$ be an $n \times 1$ vector where S is a selection matrix that picks the variables which are observable and interesting from the point of view of the analysis.

Let x_{it} be a vector of size $n \times 1$ of an observable time series filtered with method i = $1, 2, \ldots, q$ and let $x_t = [x'_{1t}, x'_{2t}, \ldots, x'_{qt}]'$. Assume that the filtered observables are linked to the true cyclical component with the structure

$$x_t = \nu_0 + \nu_1 x_t^m + u_t, \quad u_t \sim (0, \Sigma_u),$$
 (11)

where $v_0 = [v_0^1, v_0^2, \dots, v_0^q]'$ is a $nq \times 1$ vector of constants, $v_1 = [v_1^1, v_1^2, \dots, v_1^q]'$ is an $nq \times n$ matrix of nonstructural parameters, v_1^i is an $n \times n$ diagonal matrix for each i, and $u_t =$ $[u'_{1t}, u'_{2t}, \dots, u'_{qt}]$ is an $nq \times 1$ vector of possibly serially correlated errors. For estimation purposes, we normalize $\nu_1^1 = I$.

Joint estimation of the structural parameters θ and the nonstructural parameters (ν_0, ν_1, Σ_u) is now possible because (10) and (11) represent a state space system with the latter being a measurement equation and the former state equations. Thus, the likelihood of (10) and (11) can be computed with the Kalman filter. If Bayesian estimation is preferred, the posterior distribution for the parameters can be obtained with Monte Carlo Markov chain simulators (see, e.g., Canova (2007)). Note that identification of x_t^m is obtained from the cross section of filters under the conditions stated in Forni et al. (2000).

In (11), different cyclical estimates x_{it} are treated as contaminated proxies of the true cyclical component x_t^m . They are contaminated because they alter the power spectrum

of the true cyclical component at some or all frequencies. The information they contain for the model relevant concepts of cyclical fluctuations is measured by ν_0 and ν_1 . Ideally, ν_0 is a vector of zeros and ν_1 is a matrix with the identity in each $n \times n$ block, so that each set of filtered data is an unbiased and perfectly correlated, although noisy, signal of the true cyclical component. In general, we expect either $\nu_0 \neq 0$ or $\nu_1^i \neq I$, $i \neq 1$, or both, for some or all i. Since $\nu_1^1 = I$, estimates of ν_1^i for $i \neq 1$ give us an idea of the amount of correlation distortions each method displays relative to the first.

While we think of (10) and (11) as a way to correct for filtering biases, one could also think of our setup as a factor model, where the model concept of cyclical fluctuations is defined as the common factor to the noisy indicators produced by the various filters; we thank a referee for suggesting such an interpretation. This idea is appealing but disregards the information that the cyclical component of a DSGE model has a particular structure.

The signal extraction setup we use is advantageous in at least two respects. First, since we do not have to arbitrarily choose one filtering approach prior to the estimation or select which shock drives the noncyclical component, we avoid specification errors. Second, our approach can use as observables the output of one-sided and two-sided filters, both of univariate and multivariate nature, and of filters which assume that cyclical and noncyclical components are correlated or not, as long as the list of filters is sufficiently rich.

We stress that our analysis is conditional on two important assumptions. First, we assume that the model that generates x_t is correctly specified; that is, there are no missing variables or shocks. When this is not the case, the interpretation of the ν 's becomes more difficult and there is no guarantee that our signal extraction approach has better properties than any of the standard approaches. Second, we assume that the cyclical and the noncyclical components are theoretically uncorrelated. While this simplifying assumption is common in the literature, the presence of a correlation among components adds misspecification and biases which are neglected in this paper.

4.1 Selecting the filters to be used in the estimation

We have mentioned that we need vectors of filtered data which are sufficiently idiosyncratic in their low frequency distortions. Hence, knowledge of features of various filters is necessary to create a list which effectively averages out the low frequency measurement errors induced by imperfect filtering.

We have also mentioned that, apart from LT, standard filters resemble high pass filters and thus tend to underestimate the low frequency contribution of the cyclical component. Therefore, in the estimation it is important to use filters which overestimate the low frequency contribution of the cyclical components. One class of filters with such a property is the cumulative operator $(1+L)^j$, $j=1,2,\ldots$. Notice that for j=1, this filter has a square gain function which is the mirror image of the FOD filter. Low pass filters can also be considered as long as the zero frequency is properly accounted for, for example, by requiring that the sum of the filter weights is zero. One can also consider Butterworth filters, where the two free parameters are chosen to let interesting frequencies (say, from 8 to 100 quarters) be passed with minor changes.

4.2 The relationship with the literature

The literature is largely silent about the issues we address in this paper. Cogley (2001) and Gorodnichenko and Ng (2010) were concerned with the problem of estimating the structural parameters of a cyclical DSGE when the trend specification is incorrect, but did not investigate what consequences imperfect filtering has on the properties of the cyclical component or their implications for structural estimates. Giannone, Reichlin, and Sala (2006) emphasized that if model variables are measured with error, the solution has a natural factor structure, and exploited this feature to compare VAR and factor model impulse responses. Rather than consider a factor structure for the endogenous variables in terms of the states, we construct an estimable structure where vectors of filtered observable data have a factor structure in terms of the variables of the model. However, as in Giannone, Reichlin, and Sala, we emphasize that low frequency measurement error may exist. Canova (2008) suggested to estimate cyclical DSGE models by specifying a flexible link between the model and the raw data; that approach is designed to deal with different sources of misspecification than those considered here. Ferroni (2009) provided a one-step approach which allows trend specifications to be tested. The paper closest in spirit to ours is Boivin and Giannoni (2005). Their main point is that the model variables do not have an exact observable counterpart and that some indicators external to the model may have important information for model variables. The idea here is somewhat similar. The cyclical component of the model does not have an exact counterpart in the data because none of the existing filters captures the time series features of the cyclical component produced by a DSGE. If different cyclical vectors have idiosyncratic error components, this error may be averaged out with our approach.

Commentators have noticed that the procedure resembles Bayesian averaging of outcomes. Two main differences set our approach apart from this procedure. First, in Bayesian averaging the weights are the posterior probabilities of each model, while here they capture the amount of information contained in the filtered data for the modelbased concept of cyclical fluctuations. Second, in Bayesian averaging the data are the same, but the models are different. Here, there is a single model, but the data used to estimate it are different. Finally, our approach has the same flavor of multivariate unobservable component filtering (see, e.g., Canova (2007)). The extraction problem applies here to vectors of filtered data rather than to a vector of raw data.

4.3 How does the procedure fare with simulated data?

To show the properties of our approach and to highlight the practical importance of appropriately choosing the list of filters, we estimate the structural parameters of the model of Section 2 using the experimental data produced by DGP2. As input in our procedure, we employ either LT, HP, and FOD filtered data (Factor 1) or HP and BP filtered data, in which case we use two smoothing constants $\lambda = 1600$ and $\lambda = 6400$ (Factor 2). As shown in Figure 1, LT, HP, and FOD filtered data display significant low frequency differences, while HP and BP filtered data have similar low frequency components. Thus, we expect a reduction of the parameter distortions in the first case, but not in the second. Since the list of filters is short, bias will not be wiped out, but improvements in the quality of the estimates could be significant.

		_	
	True	Factor 1 Median (s.e.)	Factor 2 Median (s.e.)
σ_c	1.00	1.10 (0.10)	2.18 (0.37)
σ_n	0.70	0.49 (0.05)	0.11 (0.03)
h	0.70	0.74 (0.11)	0.62 (0.02)
ϵ	7.00	6.28 (0.11)	6.27 (0.06)
$ ho_r$	0.20	0.30 (0.07)	0.28 (0.02)
$ ho_\pi$	1.30	1.46 (0.03)	1.53 (0.05)
ρ_y	0.05	0.06 (0.01)	0.32 (0.05)
ζ_p	0.80	0.85 (0.03)	0.89 (0.01)
$ ho_\chi$	0.50	0.53 (0.05)	0.65 (0.04)
ρ_z	0.80	0.66 (0.03)	0.62 (0.02)
σ_χ	1.10	0.23 (0.07)	0.21 (0.07)
σ_z	0.57	0.19 (0.04)	0.49 (0.10)
σ_v	0.12	0.09 (0.01)	0.09 (0.01)
σ_{μ}	20.64	5.21 (0.52)	4.44 (0.73)

Table 3. Posterior parameters estimates.

We employ the same Bayesian approach used in Section 2, assuming the same priors on the structural parameters and loose priors on the nonstructural parameters entering (11). In particular, we assume that each element of ν_0 is normally distributed, with mean zero and standard deviation equal to 0.5; the prior for the nonnormalized elements of ν_1 , and is normal, centered at 1, and with standard deviation 0.5; and the prior for the standard deviation of the u_t 's is inverted gamma with mean 0.0037 and standard deviation 0.0002.

Because the data set is short, we present results obtained when constants and the loadings in (11) are common across series for each filter (in this case, there are 17 nonstructural parameters). It makes sense to restrict the model this way because the distortions we emphasize are independent of the series (see, e.g., output and inflation in Table 1). In an earlier version of the paper, we also performed unrestricted estimation (which implies 32 nonstructural parameters to be estimated): the direction of the changes was similar, but the quality of the estimates worsened.

Table 3 presents the median and the standard error of the posterior for the structural parameters when all variables are filtered. Results for other specifications are provided in Appendix B. In general, the biases we noted in Table 2 are reduced with the Factor 1 specification but not with the Factor 2 specification. For example, the habit and the risk aversion parameters are better estimated, and the inflation coefficient in the Taylor rule is much closer to the true value with Factor 1. The variability of the structural shocks is still poorly estimated, but for reasons distinct from those discussed here (these parameters are weakly identified, regardless of cyclical data used). To see what Factor 1 estimates imply in terms of economically meaningful statistics, we first compute the autocorrelation function of the cyclical components of output and inflation when the posterior median estimates of the parameters are used, and compare them with the true

^aFactor 1 uses LT, HP, and FOD filtered data; Factor 2 uses HP ($\lambda = 1600$), HP ($\lambda = 6400$), and BP filtered data. The DGP features a preference shock with two components: a stationary AR(1) and a unit root. All variables are filtered prior to estimation. The sample size is T = 150.

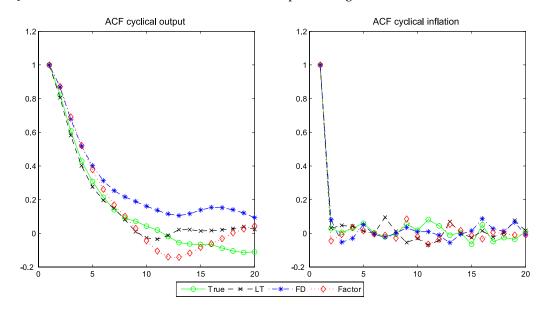


FIGURE 2. Autocorrelation functions of estimated and true cyclical components.

autocorrelation function and the autocorrelation function obtained with LT and FOD approaches (see Figure 2). For output, the autocorrelation function obtained with our specifications is very close to the true one and it is different from the one obtained, for example, with the FOD filter. For inflation, the match is good, but differences with standard methods are less dramatic, primarily because true inflation persistence is low. The good performance of our approach is reinforced when we look at the responses of the endogenous variables to the four structural shocks. Figure 3 presents the responses produced with the true parameters, those generated with the posterior median estimates obtained with our model, and those generated with LT and FOD filtered data.

Both the shape and the persistence of the conditional responses are reasonably captured by our setup. In addition, and contrary to what was happening with LT and FOD filters, the real wage response to technology has the right sign on impact. Finally, our estimates roughly replicate the magnitude of the responses to both preferences and technological disturbances, while this is not the case with standard approaches.

Next, we examine the out-of-sample performance of our setup relative to traditional ones. We conduct two types of forecasting exercises. In the first, we compute the sequence of one-step-ahead forecast errors for output and inflation, when we take as parameter values the posterior median estimates and set all the shocks in the forecasting period to zero. The MSE is computed over 150 forecasting periods, with no parameter updating in the forecasting sample; the results appear in Table 4. Figure 4 traces out the one-step-ahead path of cyclical output and cyclical inflation that would obtain with posterior median estimates when monetary shocks were drawn so as to keep the nominal interest rate fixed over the forecasting path—a standard assumption in policy projections. That is, we allow the nominal interest rate to endogenously react to output and inflation, but make sure that the monetary shocks are such that the nominal rate is

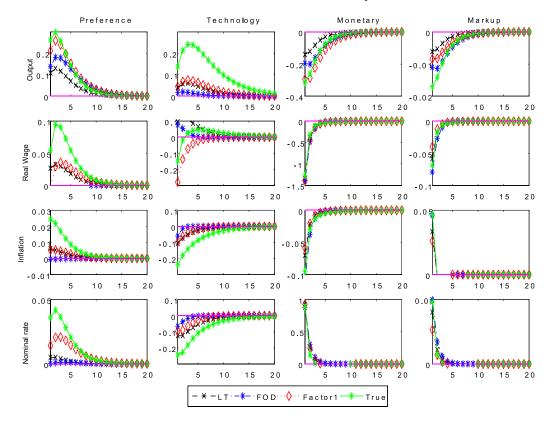


FIGURE 3. Impulse responses to shocks.

constant over the forecasting path and equal to the value taken prior to the forecasting period (time 0 in the Figure 4).

Overall, our specification is superior to single filtering approaches in unconditionally forecasting one-step-ahead cyclical output and cyclical inflation, and for output, the reduction in MSE is considerable. Our specification does well also in conditional forecasting. The counterfactual path for output that our specification produces is very close to the true one at all horizons, and practically eliminates the systematic bias that LT and FOD filters generate. For inflation, the counterfactual path produced by our model is similar to the true path; it is significantly better than the path obtained with FOD estimates, but roughly comparable to the one produced by LT estimates.

Since these conclusions hold also for alternative DGPs and combinations of filtered and unfiltered observables, the specification is effective in reducing low frequency mea-

Table 4. Mean square error of the unconditional forecasts: simulated data; scale 10^{-2} .

Series	LT	FOD	Factor 1
Output	0.006	0.003	0.001
Inflation	0.030	0.031	0.029

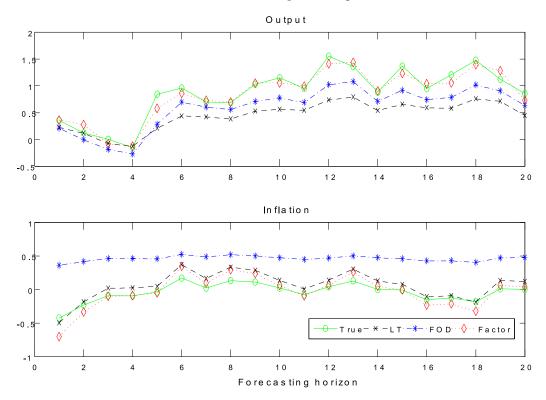


FIGURE 4. One-step-ahead forecasts, conditional on a constant interest rate path.

surement errors and can provide a more reliable picture of the cyclicality of the variables of interest.

5. Does money matter in transmitting monetary business cycles?

To show that our procedure may be relevant for understanding important economic phenomena, we reconsider the role of money in transmitting monetary business cycles. The majority of the monetary models nowadays used in the policy and academic literature attribute a minimal role to the stock of money. In most cases, these models make no reference whatsoever to monetary aggregates, and when they do, they use a specification where a money demand function determines how much money needs to be supplied, given predetermined levels of output, inflation, and the nominal rate. Ireland (2004) constructed a specification in the class of new-Keynesian models where real balances may influence the dynamics of output and inflation. He estimated the relevant parameters by likelihood techniques using post 1980 U.S. data and found that current theoretical practices are, by and large, appropriate. To construct the likelihood of his cyclical model, Ireland took away a linear trend from per capita gross domestic product (GDP) and per capita real balances, and demeaned inflation and the nominal interest rate. Here, we repeat Ireland's exercise using a number of filtering procedures.

5.1 The model economy

Since the economy is quite standard, we only briefly describe its features. At each t, the representative household maximizes

$$E_t \sum_{t} \beta^t \chi_t \left[U\left(c_t, \frac{M_t}{p_t e_t}\right) - \eta n_t \right], \tag{12}$$

where $0 < \beta < 1$, $\eta > 0$, subject to the sequence of budget constraints

$$M_{t-1} + \operatorname{Tr}_t + B_{t-1} + W_t n_t + D_t = p_t c_t + \frac{B_t}{R_t} + M_t,$$
(13)

where c_t is consumption, n_t are hours worked, p_t is the price level, M_t are nominal balances, W_t is the nominal wage, and B_t are one period nominal bonds with gross nominal interest rate R_t ; Tr_t are lump sum nominal transfers made by the monetary authority at the beginning of each t, and D_t are nominal dividends distributed by the intermediate firms; χ_t and e_t are disturbances to preferences and the money demand whose properties are described below. Let $m_t \equiv \frac{M_t}{p_t}$ denote real balances and let $\pi_t \equiv \frac{p_t}{p_{t-1}}$ denote the period t gross inflation rate.

The representative final good producing firm uses y_t^i units of intermediate good i, purchased at the price p_t^i , to manufacture y_t units of final goods according to the constant returns to scale technology $y_t = [\int_0^1 (y_t^i)^{(\epsilon-1)/\epsilon} di]^{\epsilon/(1-\epsilon)}$, where $\epsilon > 1$ is the constant price elasticity of demand for each intermediate good. Profit maximization produces the demand functions

$$y_t^i = \left(\frac{p_t^i}{p_t}\right)^{-\epsilon} y_t. \tag{14}$$

Competition within the sector implies that $p_t = (\int_0^1 (p_t^i)^{1-\epsilon} di)^{1/(1-\epsilon)}$.

The intermediate good producing firm $i \in [0,1]$ hires n_t^i units of labor from the representative household to produce y_t^i units of intermediate good i using the production function $y_t^i = z_t n_t^i$, where z_t is an aggregate productivity shock. Since intermediate goods substitute imperfectly for one another in producing finished goods, intermediate firms can set the price of their good but must satisfy (14) at the chosen price. We assume a quadratic cost in adjusting prices—measured in finished goods—given by $\frac{\phi}{2}(\frac{p_t^i}{\pi^s p_{t-1}^i} - 1)^2 y_t, \text{ where } \phi > 0 \text{ and } \pi^s \text{ measures steady state inflation. Optimal prices are chosen to maximize}$

$$E\sum_{t} \beta^{t} \chi_{t} \left[U_{1}\left(c_{t}, \frac{M_{t}}{p_{t}e_{t}}\right) \right] \left(\frac{D_{t}^{i}}{p_{t}}\right) \tag{15}$$

subject to (14), where $\beta^t \chi_t U_1(c_t, \frac{M_t}{p_t e_t})$ measures the marginal value to the household of an additional unit of profit t and real dividends are $\frac{D_t^i}{p_t} = (\frac{p_t^i}{p_t})^{1-\epsilon} y_t - (\frac{p_t^i}{p_t})^{-\epsilon} (\frac{w_t y_t}{z_t}) - \frac{\phi}{2} (\frac{p_t^i}{\pi p_{t-1}^i} - 1)^2 y_t$.

The monetary authority sets the nominal interest rate according to

$$R_t = R_{t-1}^{\rho_r} y_{t-1}^{(1-\rho_r)\rho_y} \pi_{t-1}^{(1-\rho_r)\rho_{\pi}} \Delta M_t^{(1-\rho_r)\rho_m} v_t, \tag{16}$$

where ρ_r , ρ_y , ρ_m , $\rho_m \ge 0$ are parameters and v_t is a monetary policy shock.

The law of motion of the disturbances $d_t = (\chi_t, e_t, z_t, v_t)$ is $\log d_t = \bar{d} + H \log d_{t-1} + \iota_t$, where H is diagonal with entries ρ_{χ} , ρ_e , ρ_z , and 0, respectively. The covariance matrix of the structural shocks Σ is diagonal with entries σ_{χ}^2 , σ_e^2 , σ_z^2 , and σ_v^2 . In a symmetric equilibrium, $y_t^i = y_t$, $n_t^i = n_t$, $p_t^i = p_t$, and $D_t^i = D_t$. Log-linearizing the model around the steady state produces the equilibrium conditions

$$\hat{y}_{t} = E_{t}\hat{y}_{t+1} - \omega_{1}((\hat{R}_{t} - E_{t}\hat{\pi}_{t+1}) - (\hat{\chi}_{t} - E_{t}\hat{\chi}_{t+1})) + \omega_{2}((\hat{m}_{t} - \hat{e}_{t}) - (E_{t}\hat{m}_{t+1} - E_{t}\hat{e}_{t+1})),$$

$$(17)$$

$$\hat{m}_t = \gamma_1 \hat{y}_y - \gamma_2 \hat{R}_t + (1 - (R^s - 1)\gamma_2)\hat{e}_t, \tag{18}$$

$$\hat{\pi}_t = \beta E_t \hat{\pi}_{t+1} + \psi \left(\frac{1}{\omega_1} \hat{y}_t - \frac{\omega_2}{\omega_1} (\hat{m}_t - \hat{e}_t) - \hat{z}_t \right), \tag{19}$$

$$\hat{R}_{t} = \rho_{r} \hat{R}_{t-1} + (1 - \rho_{r}) \rho_{y} \hat{y}_{t-1}$$

$$+ (1 - \rho_{r}) \rho_{\pi} \hat{\pi}_{t-1} + (1 - \rho_{r}) \rho_{m} (\Delta \hat{m}_{t} + \hat{\pi}) + \hat{v}_{t},$$
(20)

where

$$\omega_1 = -\frac{U_1\left(c^s, \frac{ms}{e^s}\right)}{y^s U_{11}\left(c^s, \frac{m^s}{p^s e^s}\right)},\tag{21}$$

$$\omega_2 = -\frac{m^s}{e^s} \frac{U_{12}\left(c^s, \frac{m^s}{e^s}\right)}{y^s U_{11}\left(c^s, \frac{m^s}{e^s}\right)},\tag{22}$$

$$\gamma_1 = \left(R^s - 1 + \frac{y^s r^s \omega_2}{m^s}\right) \left(\frac{\gamma_2}{\omega_1}\right),\tag{23}$$

$$\gamma_2 = \frac{R^s}{(R^s - 1)(m^s/e^s)} \left(\frac{U_2\left(c^s, \frac{m^s}{e^s}\right)}{(R^s - 1)e^s U_{12}\left(c^s, \frac{m^s}{e^s}\right) - R^s U_{22}\left(c^s, \frac{m^s}{e^s}\right)} \right), \tag{24}$$

$$\psi = \frac{\epsilon - 1}{\phi},\tag{25}$$

the superscript s denotes steady state values of the variables, U_j is the first derivative of U with respect to argument j = 1, 2, and U_{ij} is the second order derivative of U, i, j = 1, 2.

The log-linearized Euler condition (equation (17)) includes terms that involve real money balances and the money demand shocks. They drop out if and only if utility

is separable in consumption and real balances (see equation (22)). Similarly, real balances play a role in the forward looking Phillips curve (equation (19)) as long as $\omega_2 \neq 0$. Thus, real balances directly affect the determination of output and inflation if and only if real balances and consumption enter nonseparably in the utility function. On the other hand, the posited policy rule implies that the growth rate of nominal balances may influence output and inflation indirectly via interest rate determination. When $\omega_2 = \rho_m = 0$, real balances have no direct or indirect role in propagating cyclical fluctuations.

5.2 Estimation

We estimate the model with quarterly U.S. data spanning the period 1959:1–2008:2. All data come from the FRED data bank at the Federal Reserve Bank of Saint Louis and it is seasonally adjusted. For real GDP, we take the GDPC96 series, which is a chain weighted real value of domestic production, convert it into per capita terms by dividing it by the civilian noninstitutional population age 16 and over (CNP16OV), and log it. For real balances, we use the stock of money M2 (M2SL), divide it by the GDP deflator (GDPDEF), convert it into per capita terms by scaling it by the civilian noninstitutional population age 16 and over, and log it. Inflation is calculated by annualizing the quarterly growth rate of the GDP deflator and a three month T-bill (TB3M) is our measure of interest rates.

We employ eight procedures to extract the cyclical component of all variables. The first (POLY) fits a second order polynomial to each series separately, allowing for a change in the parameters at 1980:3. The cyclical component is the residual of the regression. The second transformation takes the first difference of each series (FOD) as an estimate of the cyclical component. The third and the fourth transformations are obtained with a HP filter and $\lambda = 1600$ or $\lambda = 128,000$; the latter leaves cycles with 2– 100 quarters of periodicity almost unchanged. The fifth transformation takes the first cumulant of all series as an estimate of the cyclical component (CUM). The sixth transformation is a multivariate version of the Beveridge and Nelson decomposition (MBN) which fits a VAR with six lags to the growth rate of the four variables and, as an estimate of the cyclical component, takes the difference between the level of the variables and their estimated long run values. The seventh transformation is a classical decomposition (CD) which assumes an additive representation of the components, fits a linear trend to the log data, and takes the residuals as the cyclical component. The last transformation employs an unobservable component (UC) decomposition which assumes that the noncyclical component is a random walk and that the cyclical component has a trigonometric representation (see Canova (2007)). Since each series has an autoregressive integrated moving average (ARIMA) (2, 1, 0) representation, the cyclical component is estimated with the projected values of an AR(2) regression of the growth rate of each variable.

We have selected these procedures to introduce as much cross-sectional idiosyncrasy in the vectors of observables as possible. In fact, in some procedures the noncyclical component is quasi-deterministic (CD, POLY), in some it is very volatile (FOD, UC, MBN), and in some it is stochastic but smooth (HP); most decompositions use univariate and one (MBN) uses multivariate information; most imply that cyclical and noncyclical components are independent and one implies that they are correlated (MBN).

Finally, some are two-sided (such as the HP filters) and some are one-sided (such as the MBN or UC filters). Note that as far as low frequency distortions are concerned, CD, POLY, CUM, and HP128,000 are likely to overestimate the low frequency variability of the cyclical component, while the other four are likely to underestimate it.

We estimate the parameters of the model by Bayesian methods. The priors are given in Appendix C. The vector of observables is 32×1 (4 series, 8 filtering methods) and the vector of states is 4×1 . Since we set $\beta = 0.99$ and steady state inflation to 2 percent, there are nine structural parameters $(\omega_1, \omega_2, \psi, \gamma_1, \gamma_2, \rho_r, \rho_p, \rho_y, \rho_m) - \epsilon$ and ϕ are not separately identifiable—and seven auxiliary parameters $(\rho_{\chi}, \rho_{e}, \rho_{z}, \sigma_{\chi}, \sigma_{e}, \sigma_{z}, \sigma_{v})$ to be estimated. We parameterize the link between the model and the cyclical data with one intercept and one slope per filter, independent of the series, but we allow the idiosyncratic term to be series and filter dependent. Thus, the intercept measures the average (across series and time) bias of each procedure and the slope measures the average correlation between the data produced by each method and the relevant model-based quantities. Since we normalize the slope of the first procedure, we have a total of 47 nonstructural parameters to be estimated (8 intercepts, 7 slopes, and 32 variances). ¹

We also estimate the structural parameters of interest using Ireland's original transformation, but allow for measurement error in each of the four equations; since our approach has an idiosyncratic error built in, this is the relevant setup for comparison. For both specifications, we draw 500,000 elements of a Markov chain Monte Carlo (MCMC) chain; convergence was achieved in less than 100,000 draws, and posterior statistics are computed using one out of every 100 of the last 200,000 draws.

5.3 The results

Before presenting estimates of the relevant parameters, we briefly comment on the estimates of the nonstructural parameters we have obtained. First, the vector of ν_0 is estimated to be zero with very small standard errors—level biases appear to be absent. Since steady state information is not used in the estimation, the mean of the data may be different from the steady state of the model at the estimated parameters. The fact that this does not happen is encouraging from an estimation point of view. Second, the loadings v_1^i vary from 0.70 (with UC filtered data) to 0.86 (with CD filtered data). Thus, all filtered series are highly correlated with the respective model quantities. Finally, standard errors for each series vary across filtering methods, confirming the presence of sufficient idiosyncratic information in the vector of cyclical data we employ. Table 5 presents the marginal likelihood of the basic specification, where both the direct and the indirect effects are allowed for, and for three restricted specifications, where either the direct effect is eliminated ($\omega_2 = 0$), the indirect effect is eliminated ($\rho_m = 0$), or both are eliminated, and the estimates of ω_2 and ρ_m are obtained in the various cases. For comparison, we

¹We have also experimented with specifications which leave all the intercepts and all the slopes free or which restrict the variances of the idiosyncratic component to be either series specific (independent of the filtering method) or filter specific (independent of the series), but we discarded them because the model fit was relatively poor.

Specification	Marginal Log Likelihood	ω_2	$ ho_m$
Basic	16,274	0.44 (0.02)	0.48 (0.02)
$\omega_2 = 0$	16,237	0	0.96 (0.01)
$\rho_m = 0$	16,212	0.43 (0.02)	0
$\omega_2=0$, $\rho_m=0$	16,220	0	0
Ireland		0.03 (0.02)	0.04 (0.03)

Table 5. Marginal log likelihood and posterior estimates.

also report estimates obtained with Ireland's filtering specification. The full set of estimates is provided in Appendix C. A model where both the direct and the indirect effects of money are present is preferable in terms of in-sample fit. Furthermore, restricting both $\rho_m = 0$ and $\omega_2 = 0$ is preferable to restricting only $\rho_m = 0$. Posterior estimates confirm this conclusions: both parameters are tightly estimated, are a posteriori different from zero, and indicate that money has a moderate influence on output and inflation fluctuations. Estimates obtained with just one filter, on the other hand, imply that both the direct and the indirect effects of money are statistically small and economically unimportant. Figure 5 presents responses to unitary impulses in our basic specification and in Ireland's. Responses look qualitatively similar, but differences in the magnitude and the persistence of the responses to shocks are evident. In particular, when our approach is used, the persistence of the responses to technology shocks is reduced, and the responses to money demand shocks have different magnitude and persistence. Interestingly, both specifications produce a liquidity puzzle (expansionary monetary shocks decrease real balances rather than increasing them) and a price puzzle (expansionary monetary shocks decrease inflation rather than increasing it). We conjecture that with a more homogenous sample, say 1984-2008, both puzzles would disappear.

To sum up, in our setup money plays a role in transmitting fluctuations to output and inflation while this is not the case when a standard single filtering approach is used. Since the list of filters we have used can average out low frequency measurement errors, the conclusions obtained with our approach appear to be more credible.

6. Conclusions

This paper has three parts. In the first, we show that standard filtering methods are unable to extract the cyclical component of a DSGE model and that measurement errors distorts estimates of the structural parameters. Biases obtain because a typical cyclical DSGE model produces time series with important low frequency components. These components are treated as noncyclical by leading filtering approaches.

In the second part, we discuss how to construct a filter which takes into account the structure that a cyclical DSGE model imposes on the data. The derivation of this filter is theoretically straightforward, but it requires knowledge of the cyclical model that generates the data. Furthermore, computational complexities make its implementation on existing computers unfeasible.

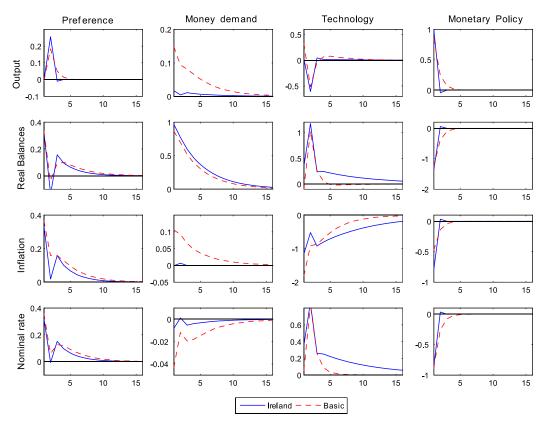


FIGURE 5. Impulse responses.

The third part proposes a method to estimate the structural parameters of a time invariant cyclical DSGE model which uses multiple sources of cyclical information. The approach borrows ideas from the recent literature that employs data-rich environments (see Boivin and Giannoni (2005)). We set up an estimation framework where the cyclical DSGE model is the unobservable factor, vectors of filtered data are contaminated observable proxies, and structural DSGE parameters are jointly estimated together with the nonstructural parameters that link the model and the observables using signal extraction techniques.

Our approach is advantageous in at least two respects. Since we do not have to arbitrarily choose one filtering method prior to estimation or select which shock drives the noncyclical component, we avoid important specification errors. Moreover, our approach can be used with cyclical data obtained with one-sided and two-sided filters, of both univariate and multivariate nature, as long as the list of filters is sufficiently rich. When appropriate conditions are satisfied, low frequency errors can be averaged out, making inference more reliable.

Using experimental data, we demonstrate that the biases obtained when just one filter is used are reduced, that the unconditional one-step-ahead mean square error (MSE) produced by our approach is smaller than the MSE obtained with a standard procedure, and that conditional forecasts are better behaved. To show that the biases are also economically relevant, we revisit the role of money in transmitting monetary business cycles. We show that when the output of multiple filters is jointly used in the estimation, money balances statistically matter for the transmission of cyclical fluctuations to output and inflation, and that the propagation of primitive shocks differs.

We want to reiterate two points which make alternatives to the procedure we present unpalatable. First, although nowadays popular, the approach of using model-based transformation to fit cyclical models is as problematic as any statistical filtering approach. Specification errors are likely to be important. Moreover, since we can solve models only when noncyclical shocks affect the technology, the consumption/investment transformation frontier, or preferences (see Chang, Doh, and Schorfheide (2007)), computational rather than economic considerations may drive model-based filtering. Thus, although some form of consistency between the model and the data is imposed, a great deal of arbitrariness is also present with this approach.

Second, the more appealing approach of employing (time varying) models to jointly explain the cyclical and the noncyclical properties of the data is currently unfeasible. Many reasons make such a research program difficult to pursue. First, jointly modelling cyclical and noncyclical fluctuations poses important theoretical challenges: there are few known mechanisms which are able to propagate temporary shocks for a long period of time (we need, for example, R&D, as in Comin and Gertler (2006), or Schumpeterian creative destruction, as in Michelacci and Lopez-Salido (2007)) or to create important cyclical implications from long run disturbances. Second, to jointly account for both types of fluctuations, we need to measure the features of noncyclical dynamics. Relatively short reliable time series and breaks of various sorts make the data largely uninformative about these features. Third, although some progress in this respect has been reported by Fernandez-Villaverde and Rubio-Ramirez (2007), time varying structures are difficult to deal with in theory and hard to handle computationally.

Given these problems, this paper provides a simple setup where specification and measurement error biases could be reduced. In this sense, the paper constitutes a step forward in improving the reliability of inferential exercises in DSGE models.

REFERENCES

Boivin, J. and M. Giannoni (2005), "DSGE estimation with data rich environments." Manuscript, University of Montreal. [75, 85, 95]

Canova, F. (1998), "Detrending and business cycle facts." *Journal of Monetary Economics*, 41, 475–512. [74]

Canova, F. (2007), *Methods for Applied Macroeconomic Research*. Princeton University Press, Princeton, New Jersey. [83, 85, 92]

Canova, F. (2008), "Bridging cyclical DSGE models and the raw data." Manuscript, UPF. [74, 85]

Canova, F. (2009), "How much structure in empirical models?" In *Palgrave Handbook of Applied Econometrics* (T. Mills and K. Patterson, eds.), Palgrave Macmillan, Basingstoke, Hampshire, U.K. [73]

Canova, F. and F. Ferroni (2011), "Supplement to 'Multiple filtering devices for the estimation of cyclical DSGE models'," Quantitative Economics, 2, http://QEconomics.org/ supp/36/supplement.pdf. [76]

Canova, F. and L. Sala (2009), "Back to square one: Identification issues in DSGE models." Journal of Monetary Economics, 56, 431–449. [74]

Chang, Y., T. Doh, and F. Schorfheide (2007), "Non stationary hours in a DSGE model." Journal of Money Credit and Banking, 39, 1357–1373. [96]

Cogley, T. (2001), "Estimating and testing rational expectations models when the trend specification is uncertain." Journal of Economic Dynamics and Control, 25, 1485–1525. [85]

Comin, D. and M. Gertler (2006), "Medium term cycles." American Economic Review, 96, 523-551. [76, 96]

Del Negro, M. and F. Schorfheide (2008), "Forming priors for DSGE models (and how it affects the assessment of nominal rigidities)." Journal of Monetary Economics, 55, 1191-1208. [74]

Fernandez-Villaverde, J. and J. Rubio-Ramirez (2007), "How structural are structural parameter values?" NBER Macroeconomic Annual, 22, 83–137. [74, 96]

Ferroni, F. (2009), "Trend agnostic one step estimation of DSGE models." MPRA Paper 14550, Library of Munich, Germany. [85]

Forni, M., M. Hallin, M. Lippi, and L. Reichlin (2000), "The generalized dynamic factor model: Identification and estimation." The Review of Economic and Statistics, 82, 540-554. [83]

Fukac, M. and A. Pagan (2010), "Limited information estimation and evaluation of DSGE models." Journal of Applied Econometrics, 25, 55-70. [74]

Gali, J. (2008), Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework. Princeton University Press, Princeton, New Jersey. [76]

Giannone, D., L. Reichlin, and L. Sala (2006), "VARs, common factors and the empirical validation of equilibrium business cycle models." Journal of Econometrics, 132, 257-279. [85]

Gorodnichenko, V. and S. Ng (2010), "Estimation of DSGE models when data are persistent." Journal of Monetary Economics, 57, 325–340. [85]

Ireland, P. (2004), "Money's role in the monetary business cycle." Journal of Money Credit and Banking, 36, 969–983. [74, 75, 89]

Justiniano, A., G. Primiceri, and A. Tambalotti (2011), "Investment shocks and the relative price of investments." Review of Economic Dynamics, 14, 101-121. [74]

Michelacci, C. and D. Lopez-Salido (2007), "Technology shocks and job flows." Review of Economic Studies, 74, 1195–1227. [76, 96]

Smets, F. and R. Wouters (2003), "An estimated dynamic stochastic general equilibrium models of the euro area." *Journal of the European Economic Association*, 1, 1123–1175. [74]

Smets, F. and R. Wouters (2007), "Shocks and frictions in US business cycles: A Bayesian DSGE approach." *American Economic Review*, 97, 586–606. [74]

Tierney, L. (1994), "Markov chains for exploring posterior distributions." *The Annals of Statistics*, 22, 1701–1762. [82]

Submitted September, 2010. Final version accepted January, 2011.