Do Seasonal Adjustments Induce Noncausal Dynamics in Inflation Rates?

Hecq, Alain and Telg, Sean and Lieb, Lenard

Maastricht University, School of Business and Economics

4 November 2016

Do Seasonal Adjustments Induce Noncausal Dynamics in Inflation Rates?

Alain Hecq **Sean Telg*** **Lenard Lieb**

Maastricht University

November 4, 2016

Abstract

This paper investigates the effect of seasonal adjustment filters on the identification of mixed causal-noncausal autoregressive (MAR) models. By means of Monte Carlo simulations, we find that standard seasonal filters might induce spurious autoregressive dynamics, a phenomenon already documented in the literature. Symmetrically, we show that those filters also generate a spurious noncausal component in the seasonally adjusted series. The presence of this spurious noncausal feature has important implications for modelling economic time series driven by expectation relationships. An empirical application on European inflation data illustrates these results. In particular, whereas several inflation rates are forecastable on seasonally adjusted series, they appear to be white noise using raw data.

Keywords: seasonality, inflation, seasonal adjustment filters, mixed causal-noncausal models.

JEL: C22, E37.

*Corresponding author: Sean Telg, Maastricht University, School of Business and Economics, P.O. Box 616, 6200 MD Maastricht, The Netherlands. E-mail: j.telg@maastrichtuniversity.nl.

We thank participants of IAAE (Milan, 2016), CFE-CM Statistics (London, 2015), ESEM-EEA (Geneva, 2016) and NESG (Leuven, 2016) for valuable comments and remarks.

1 Introduction

Most empirical macroeconomic studies are based on seasonally adjusted data. Various methods have been proposed in the literature aiming at removing unobserved seasonal patterns without affecting other properties of the time series. Just as the misspecification of a trend (at the zero frequency) may cause spurious cycles in detrended data (e.g., Nelson and Kang, 1981), a wrongly specified pattern at the seasonal frequency might have very similar effects (see, e.g., Ghysels and Perron, 1993, and references therein). In particular, (partial) autocorrelation in de-seasonalized series seems to be very prone to the chosen adjustment approach. Notably, the most popular seasonal adjustment methods used by many statistical agencies (i.e., X-11/12/13 ARIMA and TRAMO/SEATS) are based on a two-sided filter in the lag operator, and hence, impose the seasonal component to be directly related to the (partial) autocorrelation in the raw data.

Using various macroeconomic time series, Ghysels, Lee and Siklos (1993) as well as Bell and Hillmer (2002) investigate the effect of several seasonal-adjustment data transformations on serial correlation implied by univariate (causal) autoregressive moving average (ARMA) models. The estimation of ARMA models is often based on the assumption of normally distributed errors – either implicitly (OLS) or explicitly (Gaussian MLE). If this distributional assumption is supported by the data, serial correlation (and thus the autoregressive parameters) fully summarises the dependence structure of the time series. Since autocorrelation is symmetric in both calendar and reverse calendar time, no statement can be made however about the direction of causality.[1] This justifies the Box-Jenkins approach, which investigates the best fit of a model to only past values of a time series. However, if Gaussianity is violated, serial dependence may additionally involve dependence in higher order moments. This additional information proves to be useful in detecting the direction of causality; hence, the dependence with respect to the past and/or the future.

[1] The Gaussian distribution is the only distribution that is fully characterized by its autocovariance function. Since this function is symmetric (i.e. $\mathrm{Cov}(y_t, y_{t-h}) = \mathrm{Cov}(y_t, y_{t+h})$), one is unable to distinguish between forward- and backward-looking behavior.

In finance, the normality assumption is mostly abandoned, as empirical distributions of financial asset returns are often negatively skewed and more leptokurtic than a normal distribution. However, the appropriateness of Gaussianity has also been questioned for various reasons when modelling macroeconomic time series. For example, Ramsey and Rothman (1996) explain that business cycle asymmetry can be caused by the presence of non-Gaussianity in a time series. In the context of seasonal adjustment, rejecting (the assumption of) normality, potentially allows to differentiate between lagged and lead dependence of a time series. In contrast to conventional ARMA models, mixed causal-noncausal autoregressive (MAR) models formulate a process for the conditional mean in terms of serial dependence with respect to the past and/or future. That is, MAR models explicitly allow for asymmetries in lagged (causal) and lead (noncausal) dependencies. Several distributional frameworks have been proposed in the literature, such as Student's t (Lanne and Saikkonen, 2011) or LAD estimation (Hecq, Lieb and Telg, 2016). In this paper we scrutinize the impact of popular seasonal adjustment procedures on the causal/noncausal dependence structure of the adjusted series.

There is a growing consensus in the literature that in particular inflation seems to be driven by both a backward- and forward-looking component (see e.g., Galí and Gertler, 1999). MAR models constitute a natural way of modelling and forecasting such dynamics (thereby mimicking dynamics implied by various rational expectation models such as a hybrid formulation of the Phillips curve) and consequently have been entertained in various studies (see, e.g. Lanne, Luoto and Saikkonen, 2012 or Lanne and Luoto, 2012, 2013). All existing studies, however, make use of seasonally adjusted data. To what extent the noncausal dynamics are influenced by the seasonal adjustment method has not been analyzed yet. This paper aims at filling this gap in the literature.

The remainder of this paper is organized as follows. Section 2 formalizes the notion of mixed causal-noncausal models and comments on the identifiability and estimation of such models. Section 3 discusses seasonal adjustment methods based on linear filters and mentions their merits and potential pitfalls. The results of the simulation study are collected in Section 4. Section 5 details the empirical application, in which we compare MAR(r, s) model selection for both raw and seasonally adjusted quarterly inflation rates for 32 countries and one overall Europe measure. Section 6 summarizes and concludes.

2 Mixed Causal-Noncausal Models

Brockwell and Davis (1991, 2002) originally advocated the use of noncausal models as they offer the possibility to rewrite a process with explosive roots in calendar time into a process in reverse time with roots outside the unit circle. Additional important empirical features of the noncausal approach have been put forward in the recent literature. Beyond the improvement in terms of forecasting accuracy (see inter alia Lanne, Nyberg and Saarinen, 2012; Lof and Nyberg, 2015) as well as their closeness to the concept of nonfundamental shocks (see Alessi, Barigozzi and Capasso, 2011; Beaudry, Fève, Guay and Portier, 2015) simple linear noncausal models are able to mimic nonlinear processes such as bubbles or asymmetric cycles (Gouriéroux and Zakoïan, 2016; Gouriéroux and Jasiak, 2015; Hecq et al., 2016).

2.1 Model Representation

The univariate mixed causal-noncausal autoregressive model MAR(r, s) for a stationary time series $y_t, (t = 1, ..., T)$ is usually written as

$$(1 - \phi_1 L - ... - \phi_r L^r)(1 - \varphi_1 L^{-1} - ... - \varphi_s L^{-s}) y_t = \varepsilon_t, \tag{1}$$

$$\phi(L)\varphi(L^{-1}) y_t = \varepsilon_t, \tag{2}$$

with L being the backshift operator, i.e., $Ly_t = y_{t-1}$ gives lags and $L^{-1}y_t = y_{t+1}$ produces leads. When $\varphi_1 = ... = \varphi_s = 0$, the process y_t is a purely causal autoregressive process, denoted

AR(r,0) or simply AR(r):

$$\phi(L)y_t = \varepsilon_t. \tag{3}$$

Model specification (3) can be seen as the standard backward-looking AR process, with y_t being regressed on y_{t-1} up to y_{t-r}. The process in (2) becomes a purely noncausal AR$(0, s)$ model

$$\varphi(L^{-1})y_t = \varepsilon_t, \tag{4}$$

when $\phi_1 = ... = \phi_r = 0$. Model specification (4) is the counterpart of (3), since it is a purely forward-looking AR process. That is, y_t does not depend on its past values, but rather on its future values y_{t+1} up to y_{t+s}. Models of the form (2) that contain both lags and leads of the dependent variable are called mixed causal-noncausal models. In the sequel of this paper, $\phi(L)$ and $\varphi(L^{-1})$ denote the causal and noncausal polynomials, while boldfaced $\boldsymbol{\phi} = [\phi_1, ..., \phi_r]'$ and $\boldsymbol{\varphi} = [\varphi_1, ..., \varphi_s]'$ represent the corresponding parameter vectors.

The roots of both the causal and noncausal polynomials are assumed to lie outside the unit circle, that is $\phi(z) = 0$ and $\varphi(z) = 0$ for $|z| > 1$ respectively. These conditions imply that the series y_t admits a two-sided moving average (MA) representation $y_t = \sum_{j=-\infty}^{\infty} \psi_j \varepsilon_{t-j}$, such that $\psi_j = 0$ for all $j < 0$ implies a purely causal process y_t (w.r.t. ε_t) and a purely noncausal model when $\psi_j = 0$ for all $j > 0$ (Lanne and Saikkonen, 2011). In order to identify the causal from the noncausal component, the error term ε_t is assumed *iid* (and not only weak white noise) non-Gaussian.

Remark 1 *Following Gouriéroux and Jasiak (2015), we define the unobserved causal and noncausal components of the process y_t as follows:*

$$u_t \equiv \phi(L)y_t \leftrightarrow \varphi(L^{-1})u_t = \varepsilon_t, \tag{5}$$

$$v_t \equiv \varphi(L^{-1})y_t \leftrightarrow \phi(L)v_t = \varepsilon_t. \tag{6}$$

The specification of these filtered values is very useful in simulating, estimating and forecasting mixed causal-noncausal processes.

2.2 Hybrid New Keynesian Phillips Curve

A justification for the use of mixed causal-noncausal models for modelling inflation is given in Lanne and Luoto (2013). That is, they show that the hybrid New Keynesian Phillips Curve (NKPC) in its regression form,

$$\pi_t = \gamma_f \mathbb{E}_t(\pi_{t+1}) + \gamma_b \pi_{t-1} + \beta x_t + \epsilon_t,$$

where π_t denotes inflation, $\mathbb{E}_t(\cdot)$ the conditional expectation at time t, x_t a measure for marginal costs and ϵ_t an *iid* error term, can be represented as an MAR(r, s) model as in (2). Adding and subtracting $\gamma_f \pi_{t+1}$ and rearranging terms, gives

$$\pi_t = \gamma_f \pi_{t+1} + \gamma_b \pi_{t-1} + \eta_{t+1},$$

where the newly defined disturbance term η_{t+1} contains three different parts: (i) the marginal costs variable x_t, (ii) the expectation error $(\mathbb{E}_t(\pi_{t+1}) - \pi_{t+1})$ which is assumed *iid* following the literature on rational expectations models and (iii) an *iid* error ϵ_t. The newly obtained equation is divided by γ_f and lagged by one period to obtain

$$(1 - \gamma_f^{-1} L + \gamma_f^{-1} \gamma_b L^2)\pi_t = -\gamma_f^{-1} \eta_t. \tag{7}$$

Lanne and Luoto (2013) show that $a(z) \equiv (1 - \gamma_f^{-1} z + \gamma_f^{-1} \gamma_b z^2)$ can be written as the product of two polynomials, i.e., $a(z) = (1 - \phi z)(1 - \varphi^* z)$, where $\phi = \frac{1}{2}\left(\gamma_f^{-1} - \sqrt{\gamma_f^{-2} - 4\gamma_f^{-1}\gamma_b}\right)$ and $\varphi = \frac{1}{2}\left(\gamma_f^{-1} + \sqrt{\gamma_f^{-2} - 4\gamma_f^{-1}\gamma_b}\right)$ are the roots of the characteristic equation in (7). Since $|\phi| < 1$ and $|\varphi^*| > 1$ for plausible values of γ_f and γ_b (see, e.g., Galí and Gertler, 1999), the polynomial is rewritten to accommodate a noncausal representation. That is,

$$(1-\phi z)(1-\varphi^* z) = (1-\phi z)\left[-\varphi^* z(1-\frac{1}{\varphi^*}z^{-1})\right] = -\varphi^* z(1-\phi z)(1-\varphi z^{-1}),$$

with $\varphi = 1/\varphi^*$. By replacing $a(z) \equiv (1-\gamma_f^{-1}z+\gamma_f^{-1}\gamma_b z^2)$ in (7) by this newly obtained expression and allowing η_t to follow an MAR($r-1$,$s-1$) process (as η_t depends on the time series properties of x_t), one obtains:

$$(1-\phi L)(1-\varphi L^{-1})\pi_t = (\varphi^*\gamma_f)^{-1}\eta_{t+1},$$
$$\text{with } \rho(L)\theta(L^{-1})\eta_t = \zeta_t.$$

By substituting the second equation in the first one, we obtain an MAR(r,s) process like in (2) with $\phi(L) \equiv (1-\phi L)\rho(L)$, $\varphi(L) = (1-\varphi L^{-1})\theta(L^{-1})$ and a newly defined error term $\varepsilon_t = (\varphi^*\gamma_f)^{-1}\zeta_{t+1}$.

Figure 1 shows a simulated path of an MAR(2,2) with t-distributed error term. It can be seen that such models are able to capture the well-known fluctuating behavior of inflation as well as sudden peaks or throughs that could, for example, be caused by a macroeconomic policy measure. Peaks and troughs can be artifically created by considering a more leptokurtic error distribution; choosing the sum of elements in the parameter vector $\boldsymbol{\phi}$ (respectively $\boldsymbol{\varphi}$) close to unity, increases the causal polynomial (respectively noncausal) as driving force in the process.

2.3 Estimation

The non-Gaussianity assumption ensures the identifiability of the causal and the noncausal part (Breidt, Davis, Lii and Rosenblatt, 1991). Most papers by Lanne, Saikkonen and coauthors use Student's t_ν-distributions with a degree of freedom $\nu \geq 2$ as an alternative to the Gaussian distribution. Gouriéroux and coauthors rely on Cauchy or a mixture of Cauchy and Normal

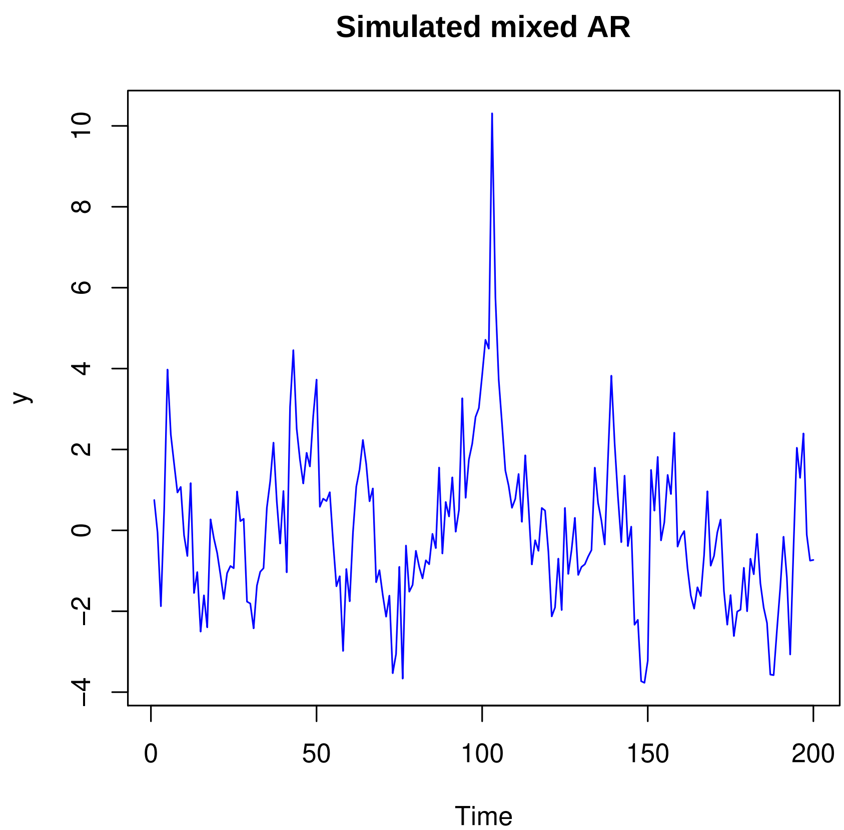

Figure 1: Simulated MAR(2,2) process with $\phi = (0.2, 0.3)$ and $\varphi = (0.3, 0.1)$, $\varepsilon_t \sim t_5$

distributions. In this paper, we also consider a non-standardized t-distribution for the error process. Lanne and Saikkonen (2011) show that the parameters of mixed causal-noncausal autoregressive models of the form (2) can be consistently estimated by approximate maximum likelihood (AML).[2] Let $(\varepsilon_1, ..., \varepsilon_T)$ be a sequence of *iid* zero mean t-distributed random variables,

[2]The term 'approximate' stems from the fact that the sample used in the likelihood contains only $T - (r + s)$ terms. As shown in Breidt et al. (1991), this quantity is only an approximation of the true joint density of the data vector $y = (y_1, ..., y_T)$.

then its joint probability density function can be characterized as

$$f_\varepsilon(\varepsilon_1, ..., \varepsilon_T|\sigma, \nu) = \prod_{t=1}^{T} \frac{\Gamma(\frac{\nu+1}{2})}{\Gamma(\frac{\nu}{2})\sqrt{\pi\nu}\sigma} \left(1 + \frac{1}{\nu}\left(\frac{\varepsilon_t}{\sigma}\right)^2\right)^{-\frac{\nu+1}{2}}.$$

The corresponding (approximate) log-likelihood function, conditional on the observed data $(y_1, ..., y_T)$ can be formulated as

$$\begin{aligned} l_y(\boldsymbol{\phi}, \boldsymbol{\varphi}, \boldsymbol{\lambda}, \alpha|y_1, ..., y_T) &= (T-p)\left[\ln(\Gamma((\nu+1)/2)) - \ln(\sqrt{\nu\pi}) - \ln(\Gamma(\nu/2)) - \ln(\sigma)\right] \\ &\quad -(\nu+1)/2 \sum_{t=r+1}^{T-s} \ln(1 + ((\phi(L)\varphi(L^{-1})y_t - \alpha)/\sigma)^2/\nu), \qquad (8) \end{aligned}$$

where $p = r + s$. The distributional parameters are collected in $\boldsymbol{\lambda} = [\sigma, \nu]'$, with σ representing the scale parameter and ν the degrees of freedom. α denotes an intercept that could be introduced in model (2), $\Gamma(\cdot)$ denotes the gamma function. Thus, the AMLE corresponds to the solution of the problem:

$$\hat{\boldsymbol{\theta}}_{ML} = \arg\max_{\boldsymbol{\theta}\in\Theta} l_y(\boldsymbol{\theta}|y_1, .., y_T),$$

with $\boldsymbol{\theta} = [\boldsymbol{\phi}', \boldsymbol{\varphi}', \boldsymbol{\lambda}', \alpha]'$ and Θ is a permissible parameter space containing the true value of $\boldsymbol{\theta}$, say $\boldsymbol{\theta}_0$, as an interior point. Since an analytical solution of the score function is not directly available, gradient based (numerical) procedures like the Berndt-Hall-Hall-Hausman (BHHH) and Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithms can be used to find $\hat{\boldsymbol{\theta}}_{ML}$. If $\nu > 2$, and hence $E(|\varepsilon_t|^2) < \infty$, the MLE is $\sqrt{T}$-consistent and asymptotically normal. Lanne and Saikkonen (2011) also show that a consistent estimator of the limiting covariance matrix is obtained from the standardized Hessian of the log-likelihood.

Hecq et al. (2016) propose an alternative method to compute standard errors which does not require the numerical optimization of the Hessian matrix and hence is more numerically stable. Denote the effective sample size $(T-p)$ used to compute the AMLE by n and define the $(1\times n)$ series $u \equiv U_t = (u_1, ..., u_{T-s})$ up to $U_{t+s} = (u_{s+1}, ..., u_T)$, $V_{t-r} = (v_1, ..., v_{T-r})$ up to $v \equiv V_t =$

$(v_{r+1}, ..., v_T)$. Then we construct the matrices $Z = [U_{t+1}, ..., U_{t+s}]'$ and $Q = [V_{t-1}, ..., V_{t-r}]'$, which are of dimensions $(s \times n)$ and $(r \times n)$ respectively. Using this notation, we can write the autoregressions defined in (5) and (6) in matrix notation as follows:

$$u = \boldsymbol{\varphi}' Z + \varepsilon, \tag{9}$$

$$v = \boldsymbol{\phi}' Q + \varepsilon. \tag{10}$$

From this it follows (see e.g., Fonseca, Ferreira and Migon, 2008) that in case of the mixed causal-noncausal model

$$\sqrt{T}(\hat{\boldsymbol{\phi}}_{ML} - \boldsymbol{\phi}_0) \xrightarrow{d} \mathcal{N}\left(0, \frac{\nu+3}{\nu+1}\sigma^2 \Upsilon_{\boldsymbol{\varphi}}^{-1}\right), \tag{11}$$

$$\sqrt{T}(\hat{\boldsymbol{\varphi}}_{ML} - \boldsymbol{\varphi}_0) \xrightarrow{d} \mathcal{N}\left(0, \frac{\nu+3}{\nu+1}\sigma^2 \Upsilon_{\boldsymbol{\phi}}^{-1}\right), \tag{12}$$

holds. We use the notation $\Upsilon_{\boldsymbol{\varphi}} = \mathbb{E}[QQ']$ and $\Upsilon_{\boldsymbol{\phi}} = \mathbb{E}[ZZ']$, where $\boldsymbol{\varphi}$ and $\boldsymbol{\phi}$ signify the relation between the unobserved values u_t, v_t and y_t as defined in (5)-(6). These quantities can be estimated consistently by $(1/n)\sum_{i=1}^{n} Q_i Q_i'$ and $(1/n)\sum_{i=1}^{n} Z_i Z_i'$, where Q_i [resp. Z_i] denotes the ith column of the matrix Q [resp. Z]. For large ν, i.e., $\nu \to \infty$, l_y approaches the Gaussian (log)-likelihood, and the model parameters cannot be consistently estimated anymore. In the empirical application, we explain in detail how to identify and select mixed causal-noncausal models for $I(0)$ series.

3 Seasonal Adjustment Methods

Seasonal adjustment of data series has received a lot of attention in the econometric and statistical literature. Many different adjustment methods have been proposed; an extensive overview can be found in Bell and Hillmer (2002). The insight that seasonality might alter the legibility of the trend and the cyclical component led to the development of moving averages that adjust series at their seasonal frequencies. These moving averages are computed using centered, symmetric linear filters of the form

$$\Psi^{SA}(L, L^{-1}) = c_0 + \sum_{j=1}^{k} c_j (L^j + L^{-j}). \tag{13}$$

As these filters are frequently used at statistical agencies, we focus on these methods in the simulation part of this paper. In particular, we focus on the linear X-11 seasonal adjustment filter, which consists of a set of moving average filters which are applied to the data sequentially (Ghysels and Perron, 1993). The linear approximation to the quarterly X-11 filter, $\Psi^{SA}_{X-11}(L, L^{-1})$, is a moving average of order 57 of which the weights sum up to one. The filter uses both leads and lags, as it takes 28 quarters before and 28 quarters after every data point into account. The final weights, rounded to 3 decimals for expository purposes, are given in Table 1.

lags/leads		lags/leads		lags/leads	
0	0.856	10	0.025	20	-0.003
1	0.051	11	0.012	21	< 0.001
2	0.041	12	-0.053	22	0.002
3	0.050	13	0.021	23	< 0.001
4	-0.140	14	0.016	24	< 0.001
5	0.055	15	-0.005	25	< 0.001
6	0.034	16	-0.010	26	< 0.001
7	0.029	17	< 0.001	27	< 0.001
8	-0.097	18	0.008	28	< 0.001
9	0.038	19	-0.002		

Table 1: Filter weights of the linear quarterly X-11 filter

The Census X-11 program was the most widely applied adjustment procedure by statistical agencies. More recent versions, such as the so-called X13-ARIMA, consist of first identifying and estimating an ARMA model on the series (with outliers, breaks, calendar effects, etc.) with the aim to extrapolate the variable in the past and in the future before taking a set of moving average filters similar to (13). This is done to preserve the number of observations that would be lost in the X-11 method without applying back- and forecasting operations. We consider a simple linear filters in this study because we want to isolate the effects coming from the moving average adjustment.

The desired property for any seasonal adjustment method is that it only affects the time series of interest at the seasonal frequencies (Ghysels et al., 1993). This requirement is not always fulfilled. For example, suppose we have a zero-mean data series y_t which is seasonally adjusted by a linear filter like in (13), with $k = 1$ (for the sake of simplicity). Then the seasonally adjusted series y_t^{SA} has first-order autocovariance equal to:

$$\begin{aligned}
\mathrm{Cov}(y_t^{SA}, y_{t-1}^{SA}) &= \mathrm{Cov}\left((\psi_1 y_{t-1} + \psi_0 y_t + \psi_{-1} y_{t+1}), (\psi_1 y_{t-2} + \psi_0 y_{t-1} + \psi_{-1} y_t)\right) \\
&= \psi_1 \psi_0 \mathbb{E}(y_{t-1}^2) + \psi_0 \psi_{-1} \mathbb{E}(y_t^2) \\
&= (\psi_0(\psi_1 + \psi_{-1}))\sigma^2 \\
&= 2\psi_0 \psi_1 \sigma^2,
\end{aligned}$$

since $\psi_i = \psi_{-i}$ for all i in the X-11 filter. That is, the seasonally adjusted series y_t^{SA} now has existing autovariances between observations at t and $t - h$ for $h \neq 0$. Since the autocovariance function is fully symmetric, these autocovariances also exist between observations t and $t + h$ for $h \neq 0$. From convolution theory, it is well-known that the autocorrelation of the linear filter is convolved with the autocorrelation of the series. In other words, the autocorrelation of $\Psi^{SA}(L, L^{-1})$ acts as a smoothing filter on the autocorrelation of y_t. However, if y_t is a white noise (or *iid*) series, these newly existing autocovariances are not due to existing dynamics in the series, but rather spuriously introduced by the seasonal adjustment filter. The special case of applying the X-11 filter to white noise has been well documented in Kaiser and Maravall

(2001).[3] We do, however, not rule out this case in our simulation study, as our main interest is to see whether seasonal adjustment filters can create spurious causal and noncausal dynamics in different data generating processes.

4 Simulation Study

We consider three data generating processes for the stationary time series $y_{1,t}, y_{2,t}$ and $y_{3,t}$

$$
\begin{aligned}
y_{1,t} &= -6D_{1,t} + 1.5D_{2,t} - 0.5D_{3,t} + 5D_{4,t} + \varepsilon_t, \\
y_{2,t} &= 0.7y_{2,t-1} + y_{1,t}, \\
y_{3,t} &= 0.7y_{3,t+1} + y_{1,t},
\end{aligned}
$$

where $D_{i,t}$ $(i = 1, ..., 4)$ are quarterly seasonal dummies with values 1 for the corresponding quarter and zero otherwise; $y_{1,t}$ is a strong white noise, $y_{2,t}$ is a causal AR(1) and $y_{3,t}$ is a noncausal AR(1). For the three processes the error term ε_t is *iid* t-distributed with 3 degrees of freedom.

On each series, we apply the X-11 linear seasonal filter $\Psi^{SA}(L)$ and perform a model selection on the adjusted series. That is, MAR(r, s) models are estimated on $y_{1,t}^{SA}$, $y_{2,t}^{SA}$ and $y_{3,t}^{SA}$ by AMLE, assuming a Student's $t-$distribution, for $r + s = p$ where $p = 1, ..., 4$ (which accounts for a total of fifteen models). We then rely on BIC for selecting the specification that minimizes that criterion. The results are collected in Table 2. We also consider model selection where, in the first step, the original variables $y_{1,t}$, $y_{2,t}$ and $y_{3,t}$ are regressed on four quarterly deterministic seasonal dummies. Model selection is performed afterwards on the residuals from this regression. The results are collected in Table 3. We display the results for three different sample sizes

[3] One of the assumptions of the X-11 method is that the series of interest y_t possesses a known ARIMA(p, d, q) structure. Subsequently, y_t is decomposed in two parts: signal s_t (the nonseasonal part) and noise n_t (the seasonal part). These parts are extracted using signal extraction theory and are assumed to follow an ARIMA process as well. In case of y_t being (strong) white noise, this assumption is flawed.

($T = 100, 400$ and 700); 1000 replications are used and we add a burn-in period of 50 observations in both sides to delete the possible effect of initial and terminal values on the simulated series.

4.1 Case 1: X-11 Seasonal Adjustment

Table 2 reports the frequencies with which BIC selects the different MAR(r, s) specifications on $y_{1,t}^{SA}$, $y_{2,t}^{SA}$ and $y_{3,t}^{SA}$. At $T = 100$, we see that the percentages with which the correct model is selected lie around 80% for all series. These results are not extremely bad, but relatively low when compared to usual results by BIC. Furthermore, it can be seen that the remaining percentages mostly go to either a MAR(r, s) of one order higher than the true data generating process or to models with $r + s = 4$, where especially the purely causal and noncausal specifications are selected. When T increases, the frequency with which a model of order 4 is selected increases by quite a margin for $y_{2,t}^{SA}$ and $y_{3,t}^{SA}$ and drastically for $y_{1,t}^{SA}$. The causal and noncausal AR(1) specifications are still selected with percentages ranging from 72% to 85%. However, these frequencies decrease as T increases (despite $T = 400$ performing better than $T = 100$). For $T = 700$, we see that in 97.3% of the cases BIC either selects a MAR(4,0) or MAR(0,4) for $y_{1,t}^{SA}$ instead of the correct white noise specification. These results are in line with Ghysels et al. (1993), who indeed find that the X-11 adjustment affects the time series properties of the data and not only at the seasonal frequencies. (Partial) autocorrelation functions are heavily affected and in our case, it seems plausible to argue that the X-11 filter creates artificial autocorrelation up to order four due to the large weight at that order in the $\Psi^{SA}(L)$ filter. In an almost equal amount of cases, the purely causal MAR(4,0) and purely noncausal MAR(0,4) maximize the log-likelihood (or similarly minimize BIC).

The effect of X-11 seasonal adjustment on a white noise series can be graphically observed in Figure 2. The graphs correspond to the spectra of the series, which (roughly said) show how much a certain frequency contributes to the series that is only observed over time. For example, peaks at frequencies $\omega = 0, \pi/2$ and π imply that the series contains a trend component and

	$T=100$			$T=400$			$T=700$		
	$y_{1,t}^{SA}$	$y_{2,t}^{SA}$	$y_{3,t}^{SA}$	$y_{1,t}^{SA}$	$y_{2,t}^{SA}$	$y_{3,t}^{SA}$	$y_{1,t}^{SA}$	$y_{2,t}^{SA}$	$y_{3,t}^{SA}$
MAR(0,0)	75.5	0.0	0.0	19.5	0.0	0.0	1.9	0.0	0.0
MAR(1,0)	6.4	82.1	4.4	4.3	83.1	0.0	0.3	73.8	0.0
MAR(0,1)	5.1	5.4	82.7	3.8	0.0	84.5	0.5	0.0	72.2
MAR(2,0)	0.6	4.1	0.2	0.2	1.8	0.0	0.0	1.9	0.0
MAR(1,1)	0.5	3.2	3.6	0.0	4.4	4.3	0.0	5.2	5.4
MAR(0,2)	0.2	0.4	3.6	0.4	0.0	1.9	0.0	0.0	1.9
MAR(3,0)	0.6	0.8	0.0	0.0	0.6	0.0	0.0	0.2	0.0
MAR(2,1)	0.1	0.3	0.3	0.0	0.2	1.0	0.0	0.0	1.2
MAR(1,2)	0.0	0.0	0.3	0.0	0.4	0.1	0.0	0.6	0.1
MAR(0,3)	0.1	0.0	0.7	0.0	0.0	0.8	0.0	0.0	0.7
MAR(4,0)	4.8	2.6	0.2	38.3	8.3	0.0	50.6	16.9	0.0
MAR(3,1)	0.0	0.2	0.8	0.2	0.0	0.9	0.0	0.0	1.8
MAR(2,2)	0.7	0.3	0.3	0.0	0.0	0.1	0.0	0.1	0.0
MAR(1,3)	0.3	0.3	0.1	0.0	1.2	0.0	0.0	1.3	0.0
MAR(0,4)	5.1	0.3	2.8	33.3	0.0	6.4	46.7	0.0	16.7

Table 2: Frequency (in percentages) with which model is selected (X-11 seasonal adjustment)

a seasonal component, associated with the once- and twice-a-year frequencies. This is a well-known pattern in quarterly time series data. As we can see in Figure 2, the original white noise series has a flat spectrum (by definition). For the X-11 seasonally adjusted white noise series, this is no longer the case. The second graph shows that the dynamics of the series are affected in such a way that certain frequencies now contribute more to the series than other. In fact, the spectrum even coincides more with that of a MAR(4,0) which is the third picture depicted in this figure. Kaiser and Maravall (2001) comment on this feature by stating that the X-11 procedure is likely to induce spurious cycles in a white noise series.

For the first-order autoregressive processes, we see a different pattern arising. Since computing spectra is fully based on the autocovariance generating function, one is unable to distinguish between purely causal and noncausal specifications. This means that the spectra of an MAR(1,0) and MAR(0,1) look exactly the same. For this reason, Figure 3 only shows the spectra of the unadjusted and adjusted MAR(1,0) process. Obviously, we observe that the X-11 seasonal ad-

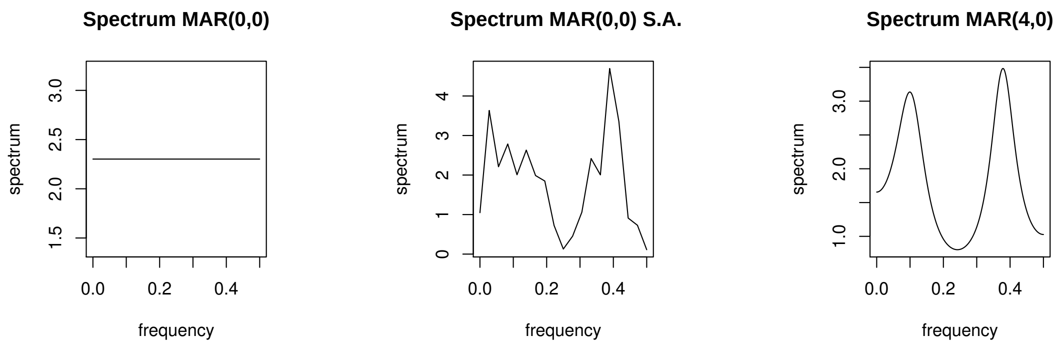

Figure 2: Spectra of (a) white noise, (b) seasonally adjusted white noise and (c) an MAR(4,0).

justment filter has affected the spectrum of the series, however only moderately when compared to the white noise case. This is in line with the results of the simulation study, where in most cases the correct model is still selected. However, the time series properties are affected in such a way that higher order models (in particular of order 4) are substantially overselected by BIC, despite the fact that BIC is known for selecting parsimonious models. Interestingly, information in other than first and second order moments seems to remain intact, since BIC selects almost only causal [noncausal] models for the seasonally adjusted causal [noncausal] DGP.[4]

4.2 Case 2: Deterministic Seasonal Adjustment

If we now use seasonal dummy adjusted variables instead of the seasonal filtered ones, we do not see the patterns of Case 1. It can be seen in Table 3 that the amount of times the correct model is selected is high. Frequencies increase with the number of time observations, making the selection consistent. In the few cases the right model is not selected, the chosen model has at most a single order more than the correct specification. For $y_{1,t}$ it selects almost equally the MAR(1,0) and MAR(0,1), while for $y_{2,t}$ it picks either the causal MAR(2,0) or the MAR(1,1).

[4] The same simulation exercise has been done based on a seasonal adjustment method called CAMPLET (see e.g., Abeln and Jacobs, 2015), which is not based on linear filters. Results show that MAR(r, s) with $r + s = 4$ are selected in most of the cases. Causality and noncausality is mostly preserved, but not to the same extent as for the X-11. Results are available upon request.

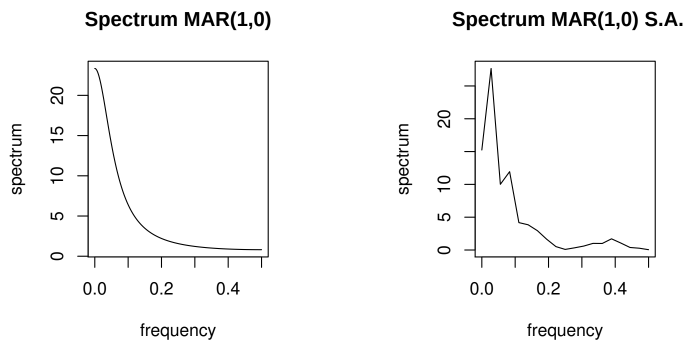

Figure 3: Spectra of (a) an MAR(1,0) and (b) a seasonally adjusted MAR(1,0).

For $y_{3,t}$, the noncausal MAR(0,2) and MAR(1,1) are often the second best choice.

	$T = 100$			$T = 400$			$T = 700$		
	$y_{1,t}$	$y_{2,t}$	$y_{3,t}$	$y_{1,t}$	$y_{2,t}$	$y_{3,t}$	$y_{1,t}$	$y_{2,t}$	$y_{3,t}$
MAR(0,0)	92.9	0.0	0.0	93.9	0.0	0.0	96.6	0.0	0.0
MAR(1,0)	2.5	89.2	2.9	2.7	95.8	0.0	1.3	96.4	0.0
MAR(0,1)	3.2	3.2	90.3	2.9	0.0	96.4	1.9	0.0	97.3
MAR(2,0)	0.3	3.6	0.1	0.2	2.1	0.0	0.1	1.5	0.0
MAR(1,1)	0.4	2.4	2.4	0.2	1.8	1.8	0.0	1.5	0.7
MAR(0,2)	0.2	0.3	2.7	0.1	0.0	1.6	0.1	0.0	1.8
MAR(3,0)	0.1	0.5	0.0	0.0	0.0	0.0	0.0	0.1	0.0
MAR(2,1)	0.1	0.2	0.5	0.0	0.3	0.1	0.0	0.2	0.0
MAR(1,2)	0.1	0.2	0.3	0.0	0.0	0.1	0.0	0.2	0.0
MAR(0,3)	0.0	0.0	0.4	0.0	0.0	0.0	0.0	0.0	0.2
MAR(4,0)	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
MAR(3,1)	0.0	0.0	0.1	0.0	0.0	0.0	0.0	0.0	0.0
MAR(2,2)	0.0	0.2	0.1	0.0	0.0	0.0	0.0	0.1	0.0
MAR(1,3)	0.0	0.2	0.1	0.0	0.0	0.0	0.0	0.0	0.0
MAR(0,4)	0.2	0.0	0.1	0.0	0.0	0.0	0.0	0.0	0.0

Table 3: Frequency (in percentages) with which model is selected (no X-11 adjustment)

In the light of cases 1 and 2, we conclude that using raw series and exploiting the correct deterministic seasonal features of the series, does not induce spurious dynamics. X-11 type of filters typically do create both causal and noncausal autoregressive parts (due to their twosidedness) even if they are applied to white noise series. We do, however, not claim that these filters should not be used. It has to be mentioned that the removal of seasonality remains a challenging task. The data generating processes considered here only take into account deterministic seasonality. In case of data containing stochastic seasonality, deterministic terms (like e.g., quarterly dummies) will not capture the true seasonal dynamics. Moreover, the power of most seasonal unit root tests is relatively low (see e.g., Del Barrio Castro, Rodrigues, Taylor, 2015), which heavily complicates the exercise of detecting their presence in the data.

5 Empirical Application

We consider Harmonized Consumer Price Index (HCPI) series for 32 European countries and one overall Europe measure (see Table 4 for countries investigated). Raw data are obtained from the Eurostat database and range from 1996Q1 until 2014Q4, which accounts for 76 quarterly observations (available for most countries). While it is obvious that prices are available monthly, we sample the monthly series (point-in-time sampling) to obtain quarterly data. Although we indeed lose time observations, we intentionally apply this transformation to compare our findings with results found in the papers quoted in the introductory section on quarterly inflation series.

We first apply a simple seasonal unit root test (HEGY test[5], see Hylleberg et al., 1990) on the natural logarithm of the raw prices, i.e.,

$$\Delta_4 y_t = \alpha + \sum_{s=1}^{3} \beta_s D_{st} + \gamma T_t + \pi_1 z_{1,t-1} + \pi_2 z_{2,t-1} + \pi_3 z_{3,t-2} + \pi_4 z_{3,t-1} + \sum_{i=1}^{p} \zeta_i \Delta_4 y_{t-i} + \varepsilon_t,$$

where D_{st} are seasonal dummies, T_t is a time trend and $\Delta_4 = (1 - L^4)$, $z_{1,t} = (1 + L + L^2 + L^3)y_t, z_{2,t} = -(1 - L + L^2 - L^3)y_t$ and $z_{3,t} = -(1 - L^2)y_t$. Three test-statistics are computed: (i) $H_0 : \pi_1 = 0$: unit root at the zero frequency (nonseasonal stochastic trend), (ii) $H_0 : \pi_2 = 0$,

[5] Alternatively, modified ($\mathcal{M}$) seasonal unit root tests, see e.g., Del Barrio Castro et al. (2015), could be used. It has been shown that these tests have good finite sample size and power properties. However, as we only apply the seasonal unit root test for illustrative purposes, we restrict ourselves to the original HEGY test.

this implies two cycles per year, (iii) $H_0 : \pi_3 = \pi_4 = 0$, the series contains roots i and $-i$ (seasonal unit roots at annual frequencies). The following transformations have to be made in order to remove the seasonal and nonseasonal unit roots in y_t: (i) if $\pi_1 = 0$, $(1 - L)$, (ii) if $\pi_2 = 0$, $(1 + L)$ and (iii) if $\pi_3 = \pi_4 = 0$, $(1 + L^2)$. The resulting transformed series is checked for an additional unit root at the zero frequency by the Augmented Dickey Fuller (ADF) test, where the standard regression equation is augmented with quarterly dummies. This additional step allows to determine the degree of integration of the inflation rate.

In the second case, the log transformed data is immediately adapted by a seasonal adjustment filter. We use the seasonal adjustment procedures in Eviews 9 (X-13 and TRAMO/SEATS) on monthly prices and afterwards we compute quarterly point-in-time sampled series. ADF tests are employed to see whether the price series are I(1) or I(2).

After having transformed prices, we apply a model selection procedure using information criteria on both raw and seasonally adjusted inflation series.[6] Since mixed causal-noncausal AR processes are not identified by Gaussian likelihood, the first step in modelling a time series with a potential forward-looking component is to check for signs of noncausality. We first estimate pseudo-causal AR(p) models by OLS and choose the model order $p = r + s$ with $p_{\max} = 8$ using BIC. Then, diagnostic tests for autocorrelation are performed to see whether additional lags are needed. The null hypothesis of normality is tested on the residuals by means of the Jarque-Bera test. In case this null cannot be rejected, there is no need to consider mixed causal-noncausal models, as the backward- and forward-looking components cannot be distinguished from each other. In case the null of normality is rejected, all MAR(r, s) specification for the selected pseudo-causal order p are considered. The model that maximizes the log-likelihood is chosen to be the final model.

[6] Matlab and $\mathcal{R}$ routines are available upon request.

5.1 Results

Table 4 shows the results of the HEGY test on the natural logarithm of raw quarterly HCPI series. Critical values are from Franses and Hobijn (1997). Rejections of the null hypotheses at a 5% significance level are indicated by asterisks. Test results indicate that the presence of a zero frequency unit root is rejected for a few countries (Czech Republic, Romania and Turkey), while the presence of seasonal unit roots at annual frequencies is rejected in all cases except for Switzerland (for which there are only 35 observations available). The possibility of prices containing two cycles per year is rejected for almost all countries, except for Lithuania, Slovenia, Spain and again Switzerland. For these countries, we report ADF results of two situations: (*i*) we ignore the seasonal unit roots and force the first difference and (*ii*) we apply the transformation $(1 + L)$ to the logarithm of the HCPI for Lithuania, Slovenia and Spain. For Switzerland, $(1 + L)(1 + L^2)$ is the correct transformation.

From the ADF test, we deduce that inflation in Bulgaria, Hungary, Ireland, Latvia, Lithuania (HEGY), Poland and Slovenia (first difference and HEGY) is not stationary. Hence, model selection is not performed for these series. We want to stress that one should not take the results of the ADF test for granted. It is known that the ADF test has relatively low power in the presence of noncausality (Saikkonen and Sandberg, 2016). Subsequently, we determine the pseudo-causal model order p, where we include an intercept and quarterly dummies in the regression equation. We find that BIC selects white noise, i.e., AR(0), for 14 of the remaining stationary series. For seven series, an AR(p) with $p > 0$ is found, but the normality of the residuals cannot be rejected. This means that we can only identify mixed causal-noncausal models on five remaining series: inflation in Greece, Iceland, The Netherlands, Slovakia and Spain (HEGY). In all cases, a model with noncausal dynamics is selected.

Table 5 shows the results for seasonally adjusted inflation data. The methods used to seasonally adjust the data are both X-13 and TRAMO/SEATS. This latter adjustment method (see Maravall, 1997), merely used by Eurostat, is based on an unobserved components decom-

	Tests on log levels			Tests on inflation			
Country	$H_0:$ $\pi_1 = 0$	$H_0:$ $\pi_2 = 0$	$H_0:$ $\pi_3 = \pi_4 = 0$	ADF statistic H_0 : unit root	Pseudo model	Jarque-Bera H_0 : normality	MAR(r, s)
Austria	-2.82	-4.51*	34.81*	-7.13*	AR(0)	reject	-
Belgium	-3.46	-4.60*	44.48*	-6.97*	AR(0)	not reject	-
Bulgaria	-0.51	-4.69*	12.31*	-2.39	-	-	-
Croatia	-1.40	-5.34*	24.96*	-6.72*	AR(0)	not reject	-
Cyprus	0.01	-4.48*	14.26*	-8.44*	AR(0)	not reject	-
Czech Republic	-3.66*	-3.94*	52.64*	-	-	-	-
Denmark	-1.89	-4.06*	50.24*	-8.06*	AR(0)	not reject	-
Estonia	-3.09	-5.44*	43.98*	-5.41*	AR(7)	not reject	-
Europe (overall)	-2.30	-5.24*	38.89*	-6.58*	AR(0)	not reject	-
Finland	-2.18	-3.39*	6.90*	-7.52*	AR(0)	reject	-
France	-2.38	-5.13*	36.08*	-7.46*	AR(0)	not reject	-
Germany	-2.48	-4.80*	33.67*	-8.11*	AR(0)	not reject	-
Greece	-2.59	-4.24*	27.77*	-3.96*	AR(4)	reject	MAR(0,4)
Hungary	-2.64	-6.60*	18.10*	-2.56	-	-	-
Iceland	-2.20	-5.05*	40.33*	-4.76*	AR(1)	reject	MAR(0,1)
Ireland	-0.44	-3.73*	65.62*	-2.66	-	-	-
Italy	-1.90	-3.48*	70.62*	-3.20*	AR(2)	not reject	-
Latvia	-1.58	-4.15*	52.58*	-2.50	-	-	-
Lithuania	-2.59	-2.70	13.57*	-2.93*	AR(6)	not reject	-
Lithuania (HEGY)	-	-	-	-2.55	-	-	-
Luxembourg	-2.27	-4.29*	34.15*	-8.11*	AR(0)	reject	-
Malta	-2.48	-4.23*	27.26*	-8.87*	AR(0)	not reject	-
Netherlands	-2.14	-3.43*	11.07*	-3.10*	AR(4)	reject	MAR(0,4)
Norway	-2.45	-6.51*	19.85*	-8.45*	AR(2)	not reject	-
Poland	-2.88	-6.19*	25.54*	-2.78	-	-	-
Portugal	-0.43	-3.32*	58.84*	-4.10*	AR(2)	not reject	-
Romania	-3.65*	-7.63*	122.90*	-	-	-	-
Slovakia	-0.73	-2.92*	64.09*	-3.17*	AR(2)	reject	MAR(0,2)
Slovenia	-1.07	-2.52	13.35*	-1.48	-	-	-
Slovenia (HEGY)	-	-	-	-2.56	-	-	-
Spain	1.00	-1.66	14.83*	-3.69*	AR(6)	not reject	-
Spain (HEGY)	-	-	-	-3.70*	AR(4)	reject	MAR(3,1)
Sweden	-1.34	-4.91*	32.04*	-8.29*	AR(0)	reject	-
Switzerland	-1.25	-0.69	4.95	-6.84*	AR(0)	not reject	-
Switzerland (HEGY)	-	-	-	-3.78*	AR(3)	not reject	-
Turkey	-4.07*	-4.66*	34.92*	-	-	-	-
United Kingdom	-1.48	-5.37*	24.87*	-7.13*	AR(0)	not reject	-
c.v. (5%)	-3.39	-2.82	6.55	-2.86		5.99	

Table 4: HEGY test on prices, ADF test and MAR(r, s) identification on quarterly inflation rates (not s.a.)

position but is not free from filters.[7] We find similar results for the two procedures. Note that we do not apply the HEGY test to the seasonally adjusted series, as seasonal effects (and thus also seasonal roots) are assumed to be removed by applying the filters. Hence, it suffices to perform the ADF-test for both the price series and inflation. Since all price series where found to contain a unit root, we only report the ADF-statistics for the rate of inflation. We find that the null hypothesis of a unit root cannot be rejected for two TRAMO/SEATS and three X-13 adjusted series. It is interesting to see that these are not always the same series; i.e., the seasonal adjustment procedure apparently directly affects the stationarity at the zero frequency (Ghysels et al., 1993).

When the pseudo-causal model is selected, the white noise specification, i.e., AR(0) with intercept, is less often selected than for the raw data. For TRAMO/SEATS only four series are found to have an AR(0) structure; for X-13 this number equals five. For the remaining pseudo-causal AR(p) models with $p > 0$, we find that the null of normality can be rejected in 15 cases for both TRAMO/SEATS and X-13. These numbers are in great contrast with the five cases that are found in the raw data application. For TRAMO/SEATS, 11 series are found to contain at least one noncausal component (eight are purely noncausal) and for X-13 this is the case for seven series (six are purely noncausal).

Hence, we find that, in general, larger order pseudo-causal models are selected by BIC when seasonally adjusted data is considered. The different methods affect the time series in such a way that different autoregressive dynamics are detected by BIC, both in the amount of AR parts and whether these are causal or noncausal. An important point to take into account is however that the number of observations for every time series is relatively low (76 at maximum).

It has to be mentioned that these findings are likely to extend to cases beyond the exercise

[7]In particular, the two-sided, centered, symmetric Wiener-Kolmogorov filter is used to estimate the signal in an observed process y_t. Theoretically, the filter uses an infinite amount of lags and leads, but in practice this is truncated to a large number m, typically representing 3-5 years of data (for more details, see e.g., Maravall, 2006).

	TRAMO/SEATS				X-13			
Country	ADF-statistic H_0 : unit root	Pseudo model	Jarque-Bera H_0 : normality	MAR(r, s)	ADF-statistic H_0 : unit root	Pseudo model	Jarque-Bera H_0 : normality	MAR(r, s)
Austria	-5.97*	AR(2)	reject	MAR(1,1)	-6.74*	AR(1)	reject	MAR(1,0)
Belgium	-5.90*	AR(1)	reject	MAR(1,0)	-6.09*	AR(1)	reject	MAR(1,0)
Bulgaria	-5.57*	AR(4)	not reject	-	-5.83*	AR(4)	not reject	-
Croatia	-7.02*	AR(1)	not reject	-	-6.48*	AR(3)	not reject	-
Cyprus	-9.61*	AR(0)	not reject	-	-9.46*	AR(0)	not reject	-
Czech Republic	-5.22*	AR(6)	not reject	-	-5.24*	AR(6)	not reject	-
Denmark	-6.87*	AR(1)	reject	MAR(0,1)	-6.84*	AR(1)	reject	MAR(0,1)
Estonia	-3.32*	AR(6)	not reject	-	-3.88*	AR(5)	not reject	-
Europe (overall)	-6.03*	AR(2)	reject	MAR(0,2)	-5.90*	AR(1)	reject	MAR(1,0)
Finland	-6.43*	AR(1)	reject	MAR(0,1)	-6.63*	AR(1)	reject	MAR(0,1)
France	-7.02*	AR(1)	not reject	-	-6.69*	AR(1)	not reject	-
Germany	-7.75*	AR(1)	not reject	-	-7.44*	AR(0)	reject	-
Greece	-4.22*	AR(2)	not reject	-	-2.88*	AR(1)	reject	MAR(0,1)
Hungary	-3.81*	AR(3)	not reject	-	-3.30*	AR(1)	not reject	-
Iceland	-5.07*	AR(1)	reject	MAR(1,0)	-2.92*	AR(2)	reject	MAR(0,2)
Ireland	-5.93*	AR(2)	not reject	-	-5.39*	AR(2)	not reject	-
Italy	-4.79*	AR(2)	reject	MAR(0,2)	-2.94*	AR(2)	reject	MAR(2,0)
Latvia	-2.87*	AR(1)	reject	MAR(0,1)	-3.14*	AR(1)	reject	MAR(1,0)
Lithuania	-4.18*	AR(2)	not reject	-	-2.83	-	-	-
Luxembourg	-7.69*	AR(0)	reject	-	-7.24*	AR(0)	reject	-
Malta	-7.69*	AR(0)	reject	-	-7.90*	AR(0)	not reject	-
Netherlands	-3.81*	AR(2)	reject	MAR(1,1)	-2.27	-	-	-
Norway	-9.04*	AR(2)	not reject	-	-9.49*	AR(2)	not reject	-
Poland	-3.29*	AR(2)	not reject	-	-3.03*	AR(2)	not reject	-
Portugal	-3.88*	AR(2)	reject	MAR(1,1)	-3.75*	AR(2)	reject	MAR(0,2)
Romania	-2.43	-	-	-	-1.88	-	-	-
Slovakia	-3.41*	AR(4)	reject	MAR(4,0)	-2.93*	AR(2)	reject	MAR(2,0)
Slovenia	-7.39*	AR(2)	reject	MAR(0,2)	-7.97*	AR(3)	reject	MAR(3,0)
Spain	-5.95*	AR(3)	reject	MAR(0,3)	-5.93*	AR(1)	not reject	-
Sweden	-7.81*	AR(0)	reject	-	-8.34*	AR(0)	not reject	-
Switzerland	-3.67*	AR(1)	reject	MAR(0,1)	-4.71*	AR(1)	reject	MAR(0,1)
Turkey	-1.82	-	-	-	-3.85*	AR(4)	reject	MAR(1,3)
United Kingdom	-6.56*	AR(1)	reject	MAR(1,0)	-6.71*	AR(4)	reject	MAR(4,0)
c.v. (5 %)	-2.86		5.99		-2.86		5.99	

Table 5: Quarterly inflation rates (s.a.) and MAR(r, s) identification

of seasonal adjustment. For instance, De Jong and Sakarya (2015) derive a new representation of the HP filter which highlights that it is a symmetric weighted average similar to the filters considered in this paper. They further state that, in case of a unit root process, the weak dependence of the cyclical component suggests that the unit root is absorbed into the trend component. As the filter (and autocorrelation function) is symmetric, this introduces spurious autocorrelation identically in calendar and reverse time.

5.2 Forecasting Inflation in Europe

We also examine the impact of seasonal adjustment filters on forecasting the inflation rate of Europe. In Table 4, an MAR(0,0) with quarterly dummies is found to be the best model for the raw data based on BIC. For the seasonally adjusted data, Table 5 indicates an MAR(0,2) and an MAR(1,0) for inflation adjusted with the TRAMO/SEATS and X-13 procedure respectively.

Figure 4 shows the raw realized rate of inflation and forecasts based on the three aforementioned models. Since the best model for the raw data does not possess an autoregressive structure, its forecasts are computed using solely deterministics. Hence, one can forecast the pattern of the quarterly dummies based on a simple OLS regression. For the other models, we find that the null hypothesis of normality can be rejected according to the results in the previous section. For this reason, we use the forecasting method proposed by Lanne et al. (2012), which can be applied directly to purely causal, purely noncausal and mixed causal-noncausal autoregressive processes. As discussed in Section 2.3, we assume that the error term follows a t_ν-distribution. From these specifications, we produce four one-step-ahead point forecasts for all seasonally adjusted series. As we are only interested in a forecasted trajectory of values, we restrict ourselves to point forecasts. Alternatively, one could consider density forecast procedures as introduced in papers by Lanne et al. (2012) and Gouriéroux and Jasiak (2015).[8]

It can be seen that no forecast is completely accurate in predicting the expected sign of the rate of inflation. Whereas the realized values indicate the presence of inflation in the second quarter and deflation in the remaining quarters, the raw forecast shows inflation in all four time periods considered. It does however mimic the movement of the realized values considerably well, as can be seen in Figures 2a and 2b. The forecasts of TRAMO/SEATS and X-13 seasonally adjusted data cannot directly be compared to the predictions of the raw data, as they are not based on the same time series process. For this reason, we cannot compute forecast performance

[8]In particular, Gouriéroux and Jasiak (2015) introduce a risk measure to assess the probability with which a bubble may burst. This method may be interesting for anticipating a large peak or trough in the rate of inflation.

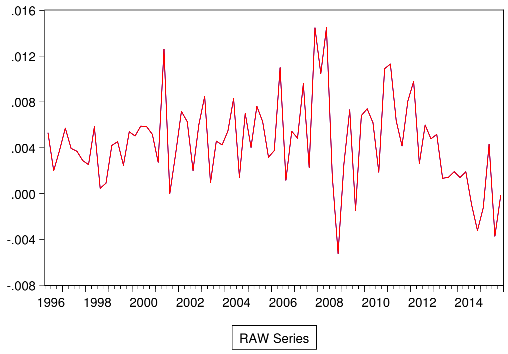

(a) Raw realized inflation (based on HCPI) from 1996Q1 until 2015Q4

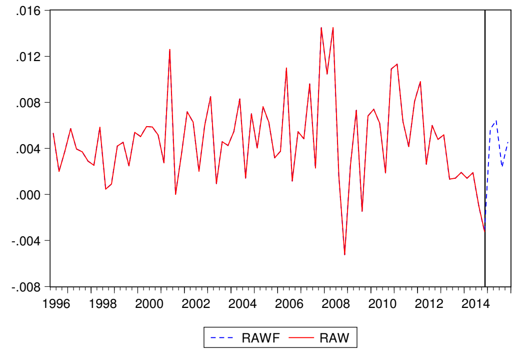

(b) Raw inflation (red), forecasts MAR(0,0) with deterministics (blue)

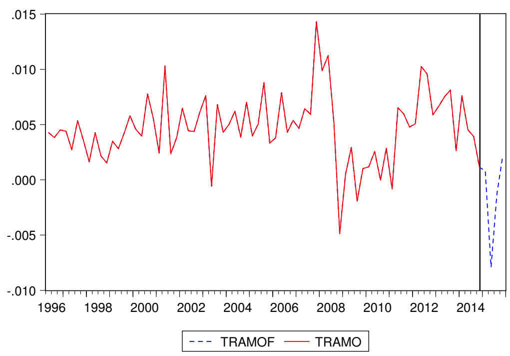

(c) TRAMO/SEATS seasonally adjusted inflation (red), forecasts (blue)

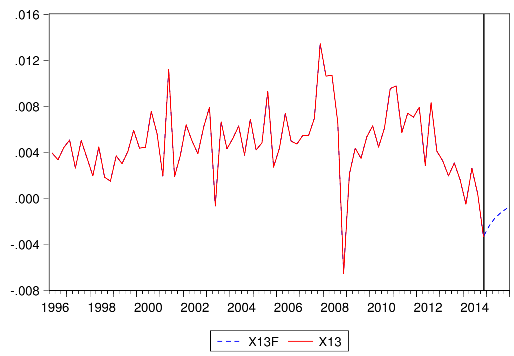

(d) X-13 seasonally adjusted inflation series (red), forecasts (blue)

Figure 4: Realized and forecasted rate of inflation for raw and seasonally adjusted data

measures such as the RMSFE. Graphically, it can be observed that the X-13 forecasts have opposite signs from the raw forecasts. The rate of deflation however rapidly becomes smaller over the four quarters, which results in the final forecasted value to come very close to the final realized value. TRAMO/SEATS overestimates both the drop in and the subsequent recovery of the inflation rate. We see that the forecasts differ between raw and seasonally adjusted data (in terms of sign and magnitude), as well as between data based on different seasonal adjustment procedures.

6 Conclusion

In this paper, we investigate the effect of seasonal adjustment on model selection for the inflation rate of 32 European countries and one overall Europe measure. In particular, we study whether seasonal adjustment may spuriously affect the noncausality found in different time series. Since raw data are directly available, we can compare model selection by BIC where one (i) deterministically removes seasonality or (ii) applies a predefined seasonal adjustment filter. We find that almost half of the series is found to be white noise in the first case, while this number is much lower in the second case. Besides, pseudo-causal models of larger order are selected in the second instance, which makes it more worthwhile to investigate the presence of noncausality (which is confirmed in approximately half of the cases). This is exactly in line with simulation results presented earlier in this paper.

As such, it seems valid to argue that model selection for mixed causal-noncausal models is heavily affected by the seasonal adjustment method performed. We do not claim that one method is better than the other, as only removing deterministic seasonality might be inappropriate when stochastic seasonality as well as breaks in seasonality are present. We do however show, by simulations, that performing seasonal adjustment on time series can create spurious autoregressive dynamics (even when the series is simply a white noise). We find that these dynamics can be both causal and noncausal.

The effects of seasonal adjustments can lead to misleading interpretations. In the empirical application, we do not only find different models for raw and seasonally adjusted data, but also find that the adjustment method matters for model selection. In a forecasting exercise on the inflation rate in Europe, we see that this leads to considerably different forecasts (both in sign and magnitude), which is a very undesirable situation for policymakers. In terms of the hybrid NKPC considered in this paper, one should be cautious when assessing the importance of forward- and backward-looking behaviour, as it might be an artefact from the filtering procedure.

References

[1] Alessi, L., Barigozzi, M. and M. Capasso (2011), Non-Fundamentalness in Structural Econometric Models: A Review, *International Statistical Review,* 79, 1.

[2] Abeln, B. and J. Jacobs (2015), Seasonal Adjustment With and Without Revisions: A Comparison of X-13ARIMA-SEATS and CAMPLET, *CAMA Working Papers 2015-25.*

[3] Beaudry, P., Fève, P., Guay, A. and F. Portier (2015), When is Nonfundamentalness in VARs a Real Problem? An Application to News Shocks, *CEPR discussion paper*, 10763.

[4] Bell, W. and S. Hillmer (2002), Issues Involved with the Seasonal Adjustment of Economic Time Series, *Journal of Business & Economic Statistics*, 20(1), Twentieth Anniversary Commemorative Issue, 98-127.

[5] Breidt, F., Davis, R., Lii, K. and M. Rosenblatt (1991), Maximum Likelihood Estimation for Noncausal Autoregressive Processes. *Journal of Multivariate Analysis*, 36, 175-198.

[6] Brockwell, P. and R. Davis (1991), *Time Series: Theory and Methods*, Springer-Verlag New York, Second Edition.

[7] Brockwell, P. and R. Davis (2002), *Introduction to Time Series and Forecasting*, Springer-Verlag New York, Second Edition.

[8] De Jong, R. and N. Sakarya (2015), The Econometrics of the Hodrick-Prescott Filter, *Working Paper*.

[9] Del Barrio Castro, T., Rodrigues P. and R. Taylor (2015), Semi-Parametric Seasonal Unit Root Tests, *DEA Working Papers*, 72.

[10] Fonseca, T., Ferreira, M. and H. Migon (2008), Objective Bayesian Analysis for the Student-t Regression Model, *Biometrika*, 95, 325-333.

[11] Franses, Ph. H. and B. Hobijn (1997), Critical Values for Unit Root Tests in Seasonal Time Series, *Journal of Applied Statistics*, 24, 25-47.

[12] Galí, J. and M. Gertler (1999), Inflation Dynamics: a Structural Econometric Approach, *Journal of Monetary Economics,* 44, 195-222.

[13] Ghysels, E., Lee H. and P. Siklos (1993), On the (Mis)specification of Seasonality and its Consequences: An Empirical Investigation with U.S. Data, *Empirical Economics,* 18(4), 191-204.

[14] Ghysels, E. and P. Perron (1993), The Effect of Seasonal Adjustment Filters on Tests for a Unit Root, *Journal of Econometrics*, 55, 57-98.

[15] Gouriéroux, C. and J. Jasiak (2015), Filtering, Prediction and Simulation Methods for Noncausal Processes. *Journal of Time Series Analysis*, doi: 10.1111/jtsa.12165.

[16] Gouriéroux, C., and J.M. Zakoïan (2016), Local Explosion Modelling by Noncausal Process, *Journal of the Royal Statistical Society, Series B*, doi:10.1111/rssb.12193.

[17] Hansen, L. and T. Sargent (1991), Two Difficulties in Interpreting Vector Autoregressions, in *Rational Expectation Econometrics*, 77-119, Westview Press, Boulder.

[18] Hecq, A., Lieb, L. and S. Telg (2016), Identification of Mixed Causal-Noncausal Models in Finite Samples, forthcoming in *Annals of Economics and Statistics*.

[19] Hylleberg, S., Engle, R., Granger C. and B. Yoo (1990), Seasonal Integration and Cointegration, *Journal of Econometrics*, 44, 215-238.

[20] Kaiser, R. and A. Maravall (2001), *Measuring Business Cycles in Economic Time Series,* Springer, New York.

[21] Lanne, M. and J. Luoto (2012), Has U.S. Inflation Really Become Harder To Forecast?, *Economics Letters*, 155(3), 383-386.

[22] Lanne, M. and J. Luoto (2013), Autoregression-Based Estimation of the New Keynesian Phillips Curve, *Journal of Economic Dynamics & Control*, 37, 561-570.

[23] Lanne, M., Luoto J. and P. Saikkonen (2012), Optimal Forecasting of Noncausal Autoregressive Time Series, *International Journal of Forecasting*, 28, 623-631.

[24] Lanne, M., Nyberg, H. and E. Saarinen (2012). Does Noncausality Help in Forecasting Economic Time Series?, *Economics Bulletin*, 32(4), 2849-2859.

[25] Lanne, M. and P. Saikkonen (2011), Noncausal Autoregressions for Economic Time Series, *Journal of Time Series Econometrics*, 3(3), 1941-1928.

[26] Maravall, A. (1997), Use and (Mostly) Abuse of Time Series Techniques in Economic Analysis, *Paper Presented at the 5th CEMAPRE Conference, Lisboa*.

[27] Maravall, A. (2006), An Application of the TRAMO-SEATS Automatic Procedure; Direct Versus Indirect Adjustment, *Computational Statistics & Data Analysis*, 50(9), 2167-2190.

[28] Nelson, C. and H. Kang (1981), Spurious Periodicity in Inappropriately Detrended Time Series, *Econometrica*, 49(3), 741-751.

[29] Ramsey, J. and P. Rothman (1996), Time Irreversibility and Business Cycle Asymmetry, *Journal of Money, Credit and Banking*, 28(1), 1-21.

[30] Saikkonen, P. and R. Sandberg (2016), Testing for a Unit Root in Noncausal Autoregressive Models, *Journal of Time Series Analysis*, 37(1), 99-125.

Appendix A - Graphs

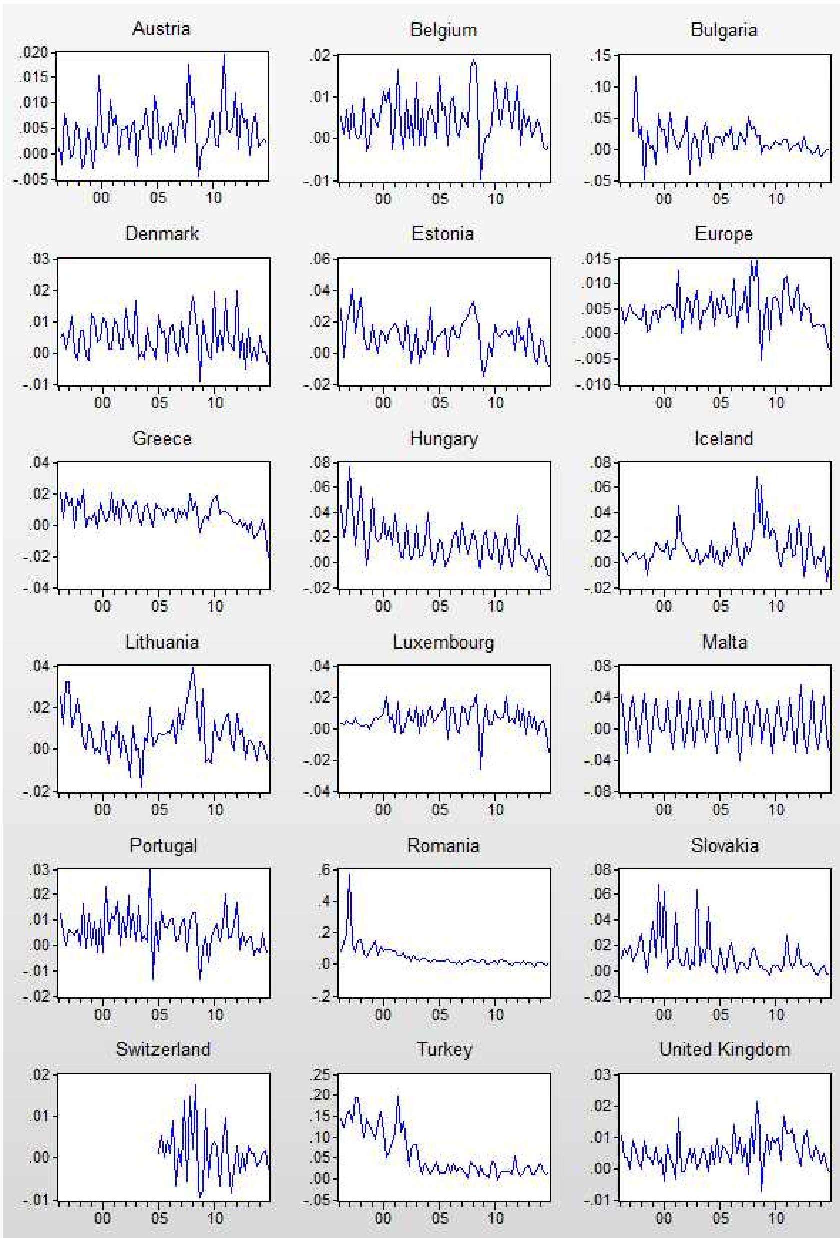

Figure 5: Raw inflation rates for first set of countries

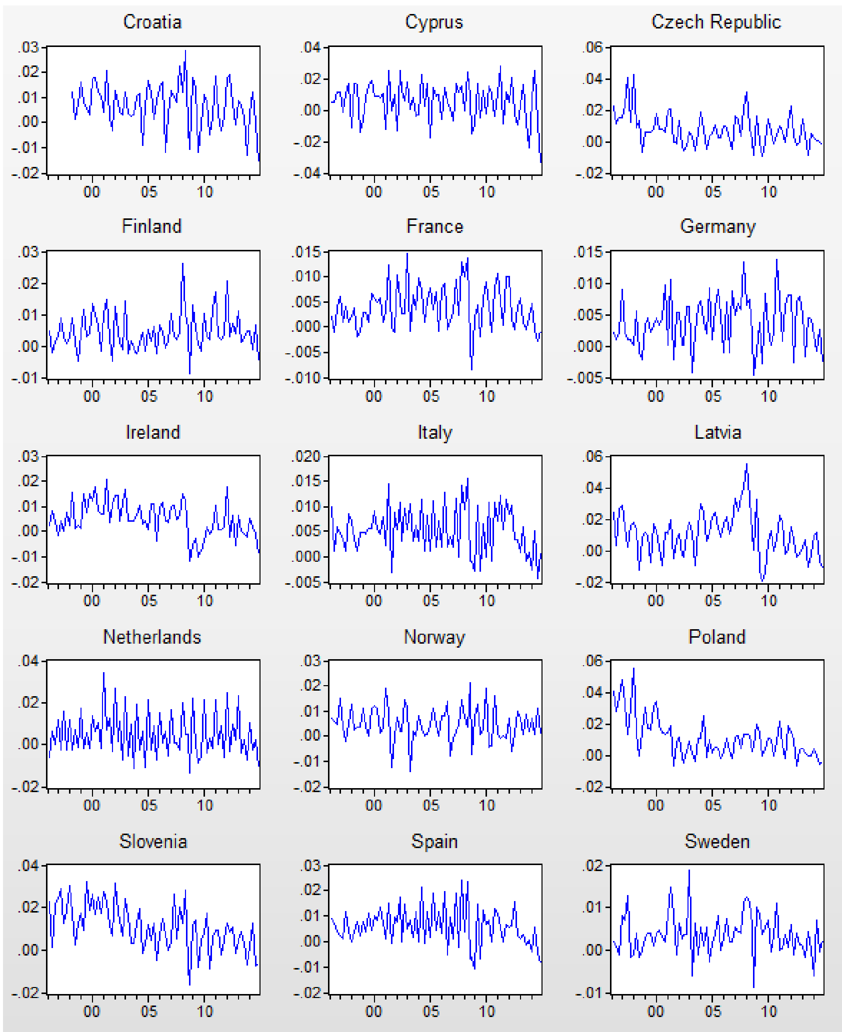

Figure 6: Raw inflation rates for second set of countries

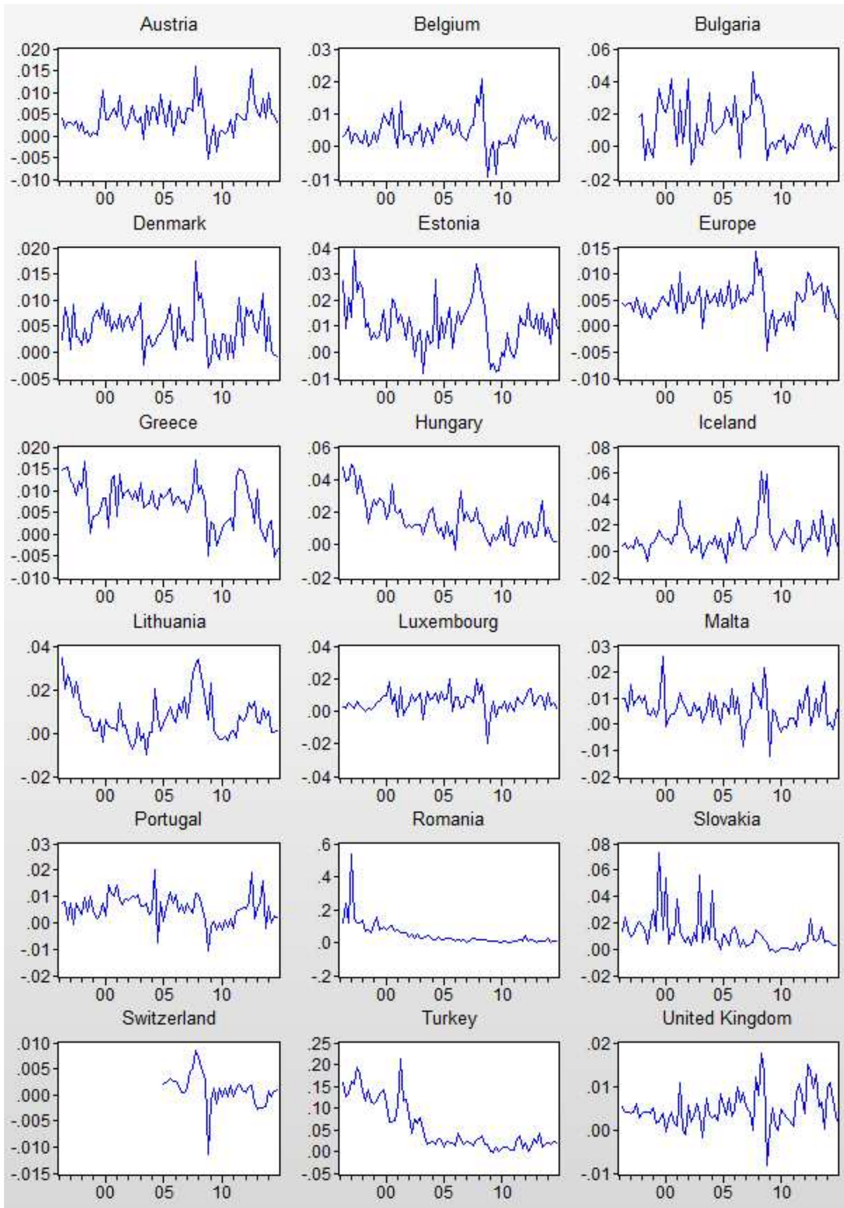

Figure 7: TRAMO/SEATS seasonally adjusted inflation rates for first set of countries

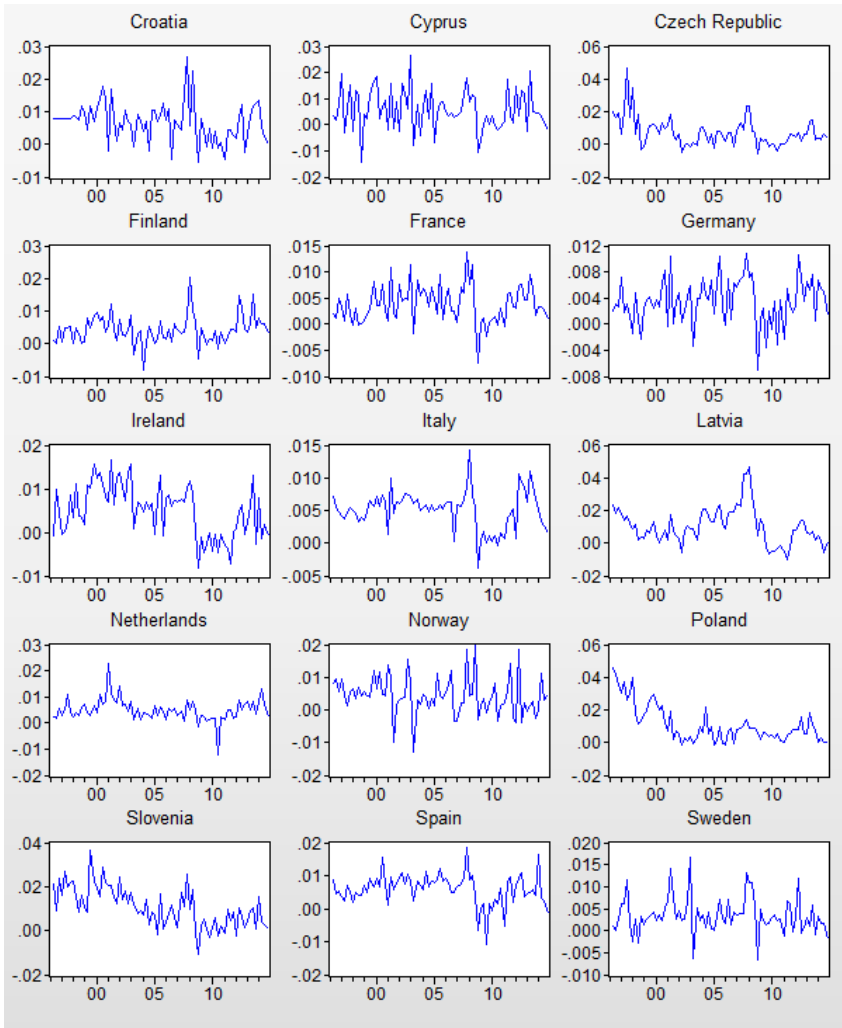

Figure 8: TRAMO/SEATS seasonally adjusted inflation rates for second set of countries