



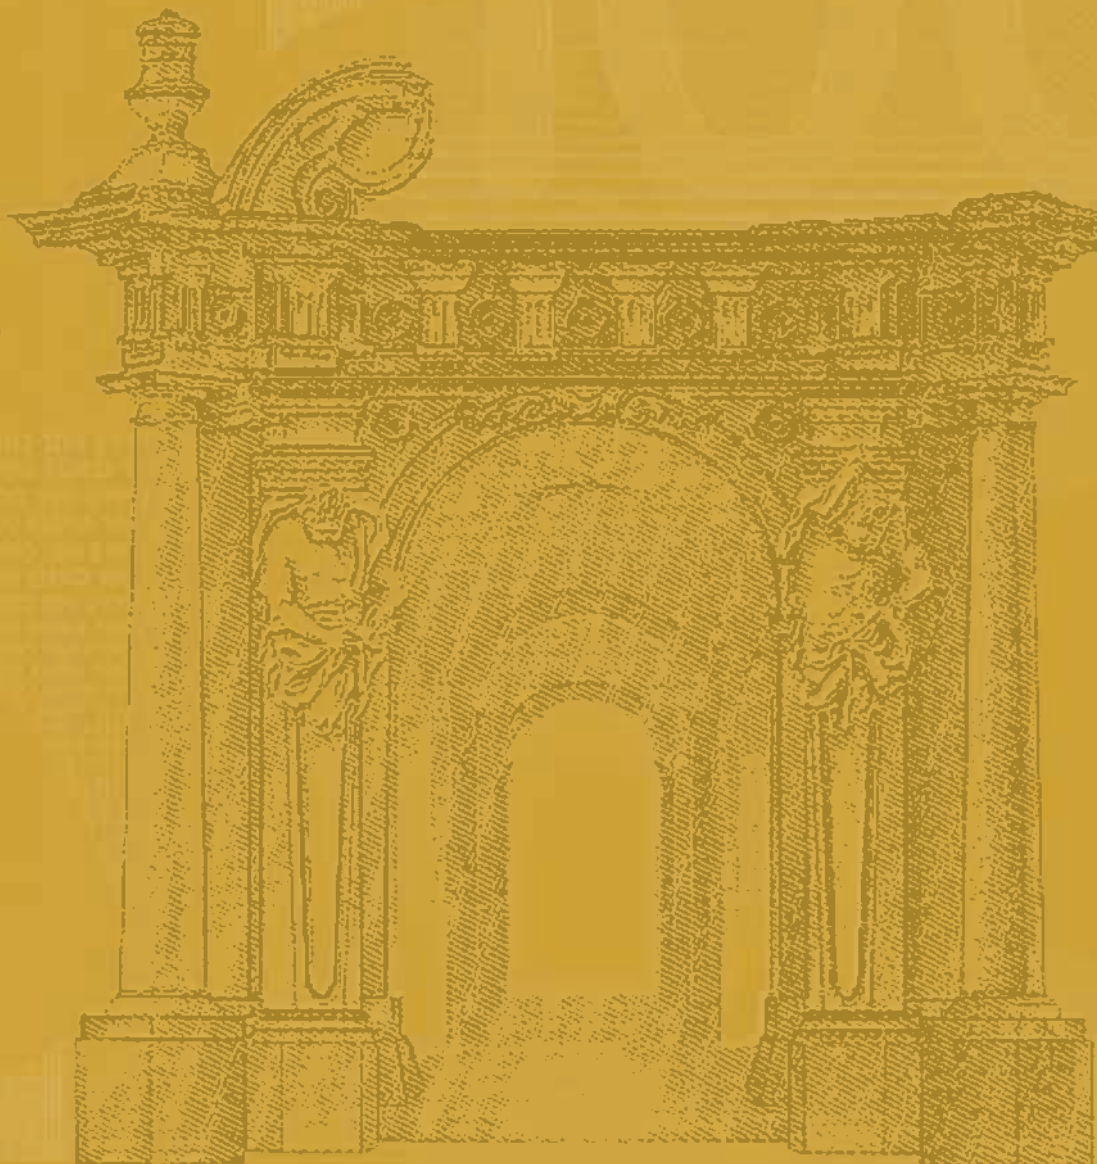
EUROPEAN CENTRAL BANK

WORKING PAPER SERIES

NO. 357 / MAY 2004

**SEASONAL
ADJUSTMENT AND
THE DETECTION
OF BUSINESS
CYCLE PHASES**

by Antonio Matas Mir
and Denise R Osborn





EUROPEAN CENTRAL BANK



WORKING PAPER SERIES

NO. 357 / MAY 2004

SEASONAL ADJUSTMENT AND THE DETECTION OF BUSINESS CYCLE PHASES ¹

by Antonio Matas Mir ²
and Denise R Osborn ³



In 2004 all publications will carry a motif taken from the €100 banknote.

This paper can be downloaded without charge from <http://www.ecb.int> or from the Social Science Research Network electronic library at http://ssrn.com/abstract_id=533026.

¹ The opinions expressed in this paper are those of the authors and do not necessarily reflect the views of the European Central Bank. The first author would like to thank Michele Manna for his support in undertaking this project. Financial support from the Economic and Social Research Council (UK) under grant number L138251030 is gratefully acknowledged by the second author.

This paper also benefited from comments and suggestions from participants at the European Meeting of the Econometric Society, Stockholm, August 2003, as well as those from an anonymous referee.

² Directorate General Statistics, European Central Bank, email: antonio.matas_mir@ecb.int.

³ Centre for Growth and Business Cycle Research, School of Economic Studies, University of Manchester, email: denise.osborn@man.ac.uk



© European Central Bank, 2004

Address

Kaiserstrasse 29
60311 Frankfurt am Main, Germany

Postal address

Postfach 16 03 19
60066 Frankfurt am Main, Germany

Telephone

+49 69 1344 0

Internet

<http://www.ecb.int>

Fax

+49 69 1344 6000

Telex

411 144 ecb d

All rights reserved.

Reproduction for educational and non-commercial purposes is permitted provided that the source is acknowledged.

The views expressed in this paper do not necessarily reflect those of the European Central Bank.

The statement of purpose for the ECB Working Paper Series is available from the ECB website, <http://www.ecb.int>.

ISSN 1561-0810 (print)

ISSN 1725-2806 (online)

CONTENTS

Abstract	4
Non-technical summary	5
1. Introduction	6
2. Seasonal adjustment filters and structural breaks	8
2.1 The X-11 filter	8
2.2 Filtering and structural breaks	11
3. Seasonal adjustment and Markov switching models: a Monte Carlo study.	16
3.1 Data generation process	17
3.1.1 DGP based on deterministic regimes	19
3.1.2 DGP based on stochastic regimes	19
3.2 Regime probabilities	20
3.2.1 Full sample smoother probabilities	20
3.2.2 Filter probabilities	21
4. Simulation results	23
4.1 Symmetric seasonal adjustment filter	23
4.1.1 DGP based on deterministic regimes	23
4.1.2 DGP based on stochastic regimes	26
4.2 One-sided seasonal adjustment filter	28
4.2.1 DGP based on deterministic regimes	28
4.2.2 DGP based on stochastic regimes	29
5. Concluding remarks	31
Appendix	32
References	35
Table and figures	38
European Central Bank working paper series	48

ABSTRACT

To date, there has been little investigation of the impact of seasonal adjustment on the detection of business cycle expansion and recession regimes. We study this question both analytically and through Monte Carlo simulations. Analytically, we view the occurrence of a single business cycle regime as a structural break that is later reversed, showing that the effect of the linear symmetric X-11 filter differs with the duration of the regime. Through the use of Markov switching models for regime identification, the simulation analysis shows that seasonal adjustment has desirable properties in clarifying the true regime when this is well underway, but it distorts regime inference around turning points, with this being especially marked after the end of recessions and when the one-sided X-11 filter is employed.

Keywords: Seasonal adjustment, business cycles, Markov switching models
JEL classifications : E32, C22, C80

Non-technical summary

The analysis of the business cycle has been following a new direction in the last ten years, with more emphasis put on the non-linear features apparently exhibited by aggregate fluctuations in real macroeconomic variables. From a theoretical perspective, structural economic models often yield non-linear time-series representations, for which linear models are only approximations. Non-linear time-series models are also useful because they allow responses that are intuitively appealing. For instance, they are capable of incorporating the possibility that economic dynamics differ in business cycle expansions and recessions. On the empirical side, however, the evidence in favour of non-linearity is still hotly debated. Nevertheless, almost all applications of non-linear models have employed seasonally adjusted data, yet little is known about the effects of seasonal adjustment on the properties of non-linear time-series models.

Standard seasonal adjustment procedures contain the implicit assumption that seasonality is uncorrelated with the business cycle. However, this is not an innocuous supposition, and may be even less so when non-linear aspects of cyclical fluctuations are to be emphasised. Against this background, we study the impact of seasonal adjustment on the detection of business cycle expansions and recessions, both analytically and empirically in the context of a Markov switching model. Analytically, we treat a business cycle regime as a structural break that is later reversed. This allows us to obtain a fairly general characterisation of the effects of seasonal adjustment on business cycle regime detection. Our results indicate that regime changes interact with the seasonal adjustment, with the interesting result that the duration of regimes influences the way in which seasonal adjustment affects the inference about the phases of the cycle. Subsequently, we study the performance of Markov switching models in tracking business cycle regimes when the data have been filtered using the X-11 US-Bureau of the Census program. Our results indicate that seasonal adjustment interferes with the tracking of the underlying regime, with the main effect being a belated detection of the end of a recession.

1. Introduction.

Since the seminal paper of Hamilton (1989), there have been a vast number of applications of regime-switching models to capture nonlinear aspects of business cycle fluctuations in key indicators such as gross domestic product or industrial production. Almost without exception these studies have employed seasonally adjusted data, which amounts to the implicit assumption that seasonality is uncorrelated with the business cycle.

This assumption has, however, been challenged by a growing body of evidence that finds an empirical link between seasonality and the business cycle (Barsky and Miron, 1989, Canova and Ghysels, 1994, Cecchetti and Kashyap, 1996, Franses and Paap, 1999, Krane and Wascher, 1999, Matas-Mir and Osborn, 2003, 2004). The clear implication of these studies is that the use of seasonally adjusted data discards information relevant for the study of business cycles. Against this, however, Christiano and Todd (2002) have recently shown that a relatively simple model of seasonality, independent of the business cycle, can capture key characteristics of the short-run dynamic relationships between observed US macroeconomic time series. Consequently, their summary view is that the use of seasonally adjusted data does not create serious distortions. Nevertheless, their data analysis is linear, and their conclusions do not necessarily extend to a nonlinear analysis of business cycle regimes.

Our purpose is to investigate further the potential interactions between seasonality and the business cycle. However, while most of the related literature to date has focussed on capturing the nature of seasonality over the business cycle, we turn our attention to the effects of conventional seasonal adjustment for the analysis of business cycle phases. Clearly, the study of seasonality and seasonal adjustment in this context are related. Nevertheless, they are also distinct, because seasonal adjustment makes specific (untested) assumptions about the nature of the seasonality in observed economic time series. Since most macroeconomic policymakers and commentators in the US and other countries rely almost exclusively on

official seasonally adjusted data, it is very important that the implications of such adjustment be understood. While there is a large literature on many aspects of the effects of seasonal adjustment, its impact on the detection of business cycle expansion and recession regimes appears to have been largely overlooked. Indeed, the only paper of which we are aware is Franses and Paap (1999), who use Monte Carlo simulations to study the differences in the estimated transition probabilities between regimes when unadjusted and adjusted data are used.

This paper examines the impact of seasonal adjustment on business cycle regime inferences. Thus, the issue of interest here is whether seasonal adjustment distorts detection of whether the economy is, for a specific period of time, in recession or expansion. The seasonal adjustment method adopted is the widely-used linear version of the X-11 program of the US Bureau of the Census, which remains the core of the new Census X-12-ARIMA program (Findley, Monsell, Bell, Otto and Chen 1998). Initially we study this issue in the context of a reversed structural break, corresponding to the economy moving from expansion to recession and then back to expansion (or vice versa). A remarkable finding is that the length of the intermediate (interpreted as recession) regime interacts with the seasonal adjustment so that the clarity of the regime switches is dependent on the duration of the intermediate regime. Subsequently we use nonlinear regime-dependent models, simulating data for both deterministic regime switches (based on the US business cycle chronology of the National Bureau of Economic Research) and stochastic switches. In both cases, a nonlinear Markov switching model is estimated. Overall, we find that seasonal adjustment acts to obscure the true underlying regime, with the primary effect being the delayed recognition of the end of recessions. In addition to the two-sided seasonal adjustment filter usually analysed, we consider the impact of the one-sided X-11 filter relevant for the analysis of the most recent data, and here the impact of seasonal adjustment is even more marked.

The plan of the paper is as follows. Section 2 gives an overview of seasonal adjustment, together with an analysis of structural breaks that provides insights into the effects of regime changes. Section 3 then explains the design of our simulation experiments for deterministic and stochastic regime switches, with results discussed in Section 4. Some conclusions complete the paper.

2. Seasonal Adjustment Filters and Structural Breaks

We first make some general points about the X-11 filter, and then (in subsection 2.2) we examine the effect of the filter on structural breaks.

2.1 The X-11 Filter

Many official statistical agencies across the world base seasonal adjustment on procedures developed within the US Bureau of the Census, specifically the X-11 program. Indeed, the X-11 seasonal adjustment filters are incorporated into many statistical software programs, allowing business and other users wide access to them. Although X-11 has recently been developed as X-12-ARIMA for use with the Bureau of the Census, the essence of this latter program relies on the same adjustment filters as in X-11 (see the discussion of X-12-ARIMA in Findley, *et al.*, 1998). Both are described in Ghysels and Osborn (2001, Chapter 5). Due to the widespread use of X-11, we concentrate our analysis on its filters. Aspects of the properties of these filters have been studied by many authors, including Bell and Hillmer (1984), Burridge and Wallis (1984), Franses and Paap (1999), Ghysels and Perron (1993, 1996), Sims (1974) and Wallis (1974). However, with the exception of Franses and Paap (1999), the effect of these filters on regime identification does not appear to have been studied previously.

We denote the original observed time series as y_t and the filtered (seasonally adjusted) series as y_t^F . Using a notation similar to Ghysels and Perron (1996), this filtering can be represented as

$$y_t^F = v(L) y_t \quad (1)$$

where, when seasonally adjusting historical data, $v(L)$ can be well approximated¹ by a symmetric two-sided linear moving average filter

$$v_2(L) = \sum_{i=-m}^m v_{2i} L^i \quad (2)$$

The coefficients v_{2i} sum to unity over $i = -m, \dots, 0, \dots, m$, with specific values for monthly and quarterly series given in Ghysels and Perron (1993). The smoothing involved in this filter at nonseasonal lags is not trivial. For example, the coefficients for the quarterly case imply positive weights of 2.1 and 1.6 percent, respectively, for observations thirteen and fourteen quarters away from observation t in both directions.

Of course, the two-sided filter (2) cannot be used to seasonally adjust the most recent observations, since it then requires unknown future observations. There are a number of possible solutions to this problem. One is to explicitly use a one-sided filter

$$v_1(L) = \sum_{i=0}^m v_{1i} L^i \quad (3)$$

(with weights again summing to unity) so that only current and past values are used; the value of m is not necessarily the same in (2) and (3). This was adopted in earlier versions of X-11, and its properties are considered by Burrige and Wallis (1984). An apparently different solution is adopted in X-12-ARIMA, as well as in X-11-ARIMA used by Statistics Canada (Dagum, 1980). These programs use a fitted ARIMA (autoregressive integrated moving

¹ Various options are available in both X-11 and X-12-ARIMA to deal with outliers, additive versus multiplicative adjustment, etc, which make the filter nonlinear; see Ghysels, Granger and Siklos (1996). Although nonlinearities introduced by seasonal adjustment are potentially important, our focus is on the impact of the linear filtering which is the core of seasonal adjustment.

average) model to generate forecasts of the required future values, and then use the two-sided filter; for example, see Findley *et al.* (1998). However, since ARIMA model forecasts are linear functions of past observations, such a procedure can be represented as a one-sided filter as in (3). The important difference in practice is that specific ARIMA models are fitted to each series, hence implying that the coefficients v_{li} differ for each series, whereas the explicitly one-sided filter approach effectively uses the same weights for all series.

Ghysels and Perron (1993) study the effect of the two-sided X-11 seasonal adjustment filter, with their analysis primarily concerned with the implications of seasonal adjustment for unit root tests. Nevertheless, they also consider the case of stationary y_t . In particular, they examine the relationship between the autocovariance at lag k for the filtered data, $\gamma^F(k)$, and the autocovariances of the unfiltered data $\gamma(j)$, $j = 0, 1, \dots$, which (assuming, without loss of generality, zero mean series) is given by

$$\begin{aligned}
 \gamma^F(k) &= E(y_t^F y_{t-k}^F) \\
 &= E \left[\sum_{i=-m_1}^{m_2} v_i y_{t-i} \sum_{j=-m_1}^{m_2} v_j y_{t-k-j} \right] \\
 &= \sum_{i=-m_1}^{m_2} \sum_{j=-m_1}^{m_2} v_i v_j E[y_{t-i} y_{t-k-j}] \\
 &= \sum_{i=-m_1}^{m_2} \sum_{j=-m_1}^{m_2} v_i v_j \gamma(k+j-i)
 \end{aligned} \tag{4}$$

This last expression can be evaluated in a straightforward way for given filter coefficients (either one-sided, with $m_1 = 0$, or symmetric two-sided, with $m_1 = m_2$) and given autocovariance properties of the original process for y_t . If (two-sided) X-11 seasonal adjustment is applied to a quarterly white noise process where $\gamma(0) = \sigma^2$ and $\gamma(j) = 0$, $j \neq 0$, then (using the weights given by Laroque, 1977, reproduced in Ghysels and Perron, 1993) the autocorrelation function of the adjusted process is shown in Figure 1. Nontrivial negative autocorrelation is induced at annual lags of one, two and three years, which can be anticipated

as the role of seasonal adjustment is to remove the strong positive autocorrelations at seasonal lags consequent upon seasonality. In addition, however, seasonal adjustment induces positive and persistent autocorrelation at nonseasonal lags.

2.2 Filtering and Structural Breaks

Our purpose is to examine the effect of seasonal adjustment on the detection of business cycle regimes. A single regime switch can be regarded as a type of structural break, with such breaks being analysed by Ghysels and Perron (1996). However, business cycle regime shifts imply that the break is later reversed, for example, when the economy moves from expansion to recession and then back to expansion. Therefore, in this subsection we extend the analysis of Ghysels and Perron to this case.

To begin the investigation with the single break case, assume that the observed series is zero prior to the break date T_B and unity thereafter, so that

$$y_t = \begin{cases} 0 & t = \dots, T_B - 2, T_B - 1 \\ 1 & t = T_B, T_B + 1, \dots \end{cases} \quad (5)$$

This assumption for y_t involves no loss of generality in the context where the structural break is a step change, since (5) can be scaled and added to a process with constant properties to create any desired change. When the two-sided filter of (2) is applied in the context of the structural break (5), the filtered series is given by

$$y_t^F = \begin{cases} 0 & t < T_B - m \\ \sum_{i=T_B-t}^m v_{2i} & T_B - m \leq t \leq T_B \\ 1 - \sum_{i=t-T_B+1}^m v_{2i} & T_B < t < T_B + m \\ 1 & t \geq T_B + m \end{cases} \quad (6)$$

where we use the symmetry of the filter and the fact that the weights sum to unity. As indicated by a comparison of (5) and (6), all filtered values for $T_B - m \leq t < T_B + m$ are



influenced by the structural break and this effect can be computed. A similar result to (6) can be obtained for the application of the one-sided filter (3), except that only observations subsequent to the break, $t \geq T_B$, are distorted.

Figure 2 illustrates the effect graphically for the two-sided quarterly X-11 filter², where the upper panel shows the filtered series, y_t^F , and the lower panel shows the distortion induced by the filter as the difference between the filtered and original series, which is (6) minus (5). A similar graph is shown by Ghysels and Perron (1996). Due to the filter, the magnitude of the break at $t = T_B$ is reduced by 7.2 percent, while the last observation in the previous regime (namely, $t = T_B - 1$) is increased by the same amount, thereby reducing the step change between these periods by more than 14 percent. Further, there is substantial distortion of the values both a year before and a year after T_B , again with distortions in consecutive quarters of ± 7 percent of the value of the break, thereby introducing spurious evidence of a step change of magnitude 14 percent one year before and one year after the true break.

In order to more accurately reflect the implications of business cycle regime changes, we now turn to the case of two off-setting structural breaks. To be more precise, assume that a unit structural break occurs at time $t = T_B$, with this break then reversed at the later time period $t = T_C$. Therefore, $y_t = 0$ except for $T_B \leq t < T_C$ when $y_t = 1$. In a business cycle context, this sequence of two breaks can be interpreted as recession followed by expansion and another recession, or (more interestingly, when scaled by -1) as expansion followed by recession and then expansion.

Defining y_{Bt} using the single break of (5) and y_{Ct} in the same way, with T_C replacing T_B , then the reversed structural break series we wish to analyse is given by

² Our analysis in this case is restricted to the two-sided filter, because the one-sided filter relates only to seasonal adjustment of the final observation of the sample. For historical data, different asymmetric filters apply to preceding observations until the two-sided symmetric filter can be applied. In order to keep the analysis simple, we do not consider this case here.

$$y_t = y_{Bt} - y_{Ct} \quad T_B < T_C \quad (7)$$

Again since the filter $v_2(L)$ is linear, it follows that the distortion from applying this two-sided filter to y_t is given by

$$y_t^F - y_t = (y_{Bt}^F - y_{Bt}) - (y_{Ct}^F - y_{Ct}) \quad (8)$$

where the superscript F indicates the corresponding filtered series. Using (6), it follows that the distortion in the series y_t of (7), in terms of the filter weights, is given by

$$y_t^F - y_t = \begin{cases} \sum_{T_B-t}^m v_{2i} & T_B - m \leq t < T_C - m \\ \sum_{T_B-t}^{T_C-t-1} v_{2i} & T_C - m \leq t < T_B \\ -2 \sum_{T_C-t}^m v_{2i} - \sum_{t-T_B+1}^{T_C-t-1} v_{2i} & T_B \leq t < T_C \\ \sum_{t-T_C+1}^{t-T_B} v_{2i} & T_C \leq t < T_B + m \\ \sum_{t-T_C+t}^m v_{2i} & T_B + m \leq t < T_C + m \\ 0 & \text{otherwise} \end{cases} \quad (9)$$

when we assume that $T_C < T_B + m$, so that the second “reversed break” at T_C occurs within the period where the filtered values are affected by the initial break at T_B . Since $m = 28$ for the quarterly X-11 filter, this is a realistic assumption in the context where the intermediate regime is recession. For $T_C \geq T_B + m$, there is little or no overlap between the values with substantial weights in the seasonal adjustment filters at the two break dates, and hence for practical purposes these can be analysed separately.

From a superficial examination of (9), it appears that the greatest distortion may be anticipated as negative effects reducing the magnitude of the change during the period of the intermediate regime, namely for $T_B \leq t < T_C$. However, this is not necessarily true, because some “weights” v_{2i} are negative. Therefore, Figure 3 plots the values given by (9) for the two-sided linear approximation to the quarterly X-11 filter for cases where $T_C = T_B + k$ for $k = 2, 3$,

4, 5, corresponding to durations of this regime between two and five quarters. Panel (a) of Figure 3 considers periods within the intermediate regime, $T_B \leq t < T_C$, while panel (b) corresponds to $t = T_C - 1 + j$ for $j = 0, \dots, 16$ and hence shows the distortion for the final period of the intermediate regime and the first 16 quarters of the new “normal” regime. Notice that, for ease of exposition, the final period of the intermediate regime is included in both panels. As implied by (9), the distortions for $T_B - j$ (prior to the first break) are symmetrical with those shown in panel (b).

When it happens that the intermediate regime lasts exactly one year, then seasonal adjustment leads to little distortion, so that both regime changes are effectively as easy to identify using the filtered or the original series. This is due to the annual cycle in the distortions induced with a single structural break (Figure 2). With two breaks, these distortions are subtracted at a lag of one year, hence virtually cancelling out. Within the intermediate regime, seasonal adjustment does generally reduce the magnitude of the structural break, although it is enhanced for the central three quarters of this regime when it lasts for five quarters.

However, the most notable effects are seen in panel (b) of Figure 3. When the regime has duration 2, 3 or 5 quarters, filtering creates very substantial distortion in relation to the magnitude of the break, with this being larger than in the case of a single structural break. More specifically, during the final period of the intermediate regime when $y_t = 1$, the magnitude of this value is reduced by more than 14 percent when the duration is five quarters and more than 5 percent of the value “leaks” to the adjacent first quarter of the regime $y_t = 0$. Thus, comparison between these values reduces the step change by almost 20 percent. The total effect is similar when the intermediate regime has durations of two or three quarters. However, in these latter cases of shorter regimes, the reduction of the end quarter of the regime $y_t = 1$ is reduced by less (5 and 9 percent respectively), while the leaking of the

intermediate regime value to the adjacent quarter of the other regime is greater. In both of these cases, the magnitude of the step change at $t = T_c$ (and at $t = T_B$) is reduced by around 20 percent.

It is also notable that the subsequent distortion follows a marked seasonal pattern. Indeed, the distortions shown in Figure 3(b) a year subsequent to the end of the intermediate regime (except for the case $k = 4$) imply a spurious break with a magnitude around 17 percent of the value of the reversed break. This is due to the negative distortion four quarters after the end of the regime combined with the positive effect in the next quarter, and could be mistaken for evidence of a return to the intermediate regime.

These effects are not simply an artefact of our assumption that the series y_t has no seasonality. Irrespective of the pattern in a series, a structural break will imply (at least for the linear approximation of the filter) that the distortion of the time series effect of the break will be unchanged from that analysed here. For example, assume that the quarterly time series y_t undergoes a shift at some time T_B . Then the series can be represented as

$$y_t = x_t + \beta b_t$$

where x_t is the time series without the break, while b_t is the zero/one break series, as given in (5) or (7), and β is a scalar reflecting the value of the break that occurs at $t = T_B$. Then since linear filtering yields

$$v(L)y_t = v(L)x_t + \beta v(L)b_t,$$

it follows that

$$y_t^F - y_t = (x_t^F - x_t) + \beta(b_t^F - b_t). \quad (10)$$

Therefore, (10) implies that the effect of filtering on y_t can be additively decomposed into the effect on the series x_t and the effect on the break series, βb_t . Put a different way, (10) implies that the seasonality attributed to y_t by the linear filter (namely, $y_t - y_t^F$) can be decomposed as

the seasonal effect obtained using the series without a structural break, $x_t - x_t^F$, plus β times the apparent “seasonal effect” induced by the break.

These effects occur because X-11 implicitly assumes evolving seasonality due to the presence of seasonal unit roots, and these are removed through the annual summation filter $1 + L + L^2 + L^3$ (assuming quarterly data). When a structural break occurs, the filter is unable to distinguish fully between changing seasonality and the nonseasonal break, and hence it effectively allocates part of the break effect to a change in the seasonal pattern, which is the effect illustrated in Figures 2 and 3.

From the point of view of identifying the period of the break, since seasonal adjustment dilutes the magnitude of the break and introduces spurious apparent shifts before and after the true break, then we might expect detection of such a break to be more difficult after seasonal adjustment. Although this has many potential implications for analysing breaks using seasonally adjusted series, our particular interest is in the ability of the Markov switching model to detect recurring shifts between regimes. It is to this issue that we now turn.

3. Seasonal Adjustment and Markov Switching Models: a Monte Carlo Study.

At first sight, we might anticipate that the smoothing of the original series analysed in the previous section will have a negative impact on the detection of business cycle regimes through the Markov switching model. However, this is not entirely clear-cut. It is true that the smoothing may partially disguise regime switches when they actually happen. However, at the same time the smoothing involved in seasonal adjustment will facilitate correct classification in periods when no regime switch takes place. For example, if the regimes capture (classical) business cycle recessions versus expansions, low growth within an expansion or zero growth within an overall contraction regime do not signal regime switches, and the smoothing effect

of X-11 may reduce the chances that these relatively extreme observations within a given regime result in an incorrect regime classification. Thus, we entertain the prior expectation that the application of X-11 to a regime-switching process may reduce the number of regime “false alarms” as well as increasing the number of regime “missed calls”. However, it is not clear what the overall effects of the combination of these two opposing forces will be.

We address these issues through Monte Carlo simulations. This section describes the models utilised for the simulations, while section 4 presents our results. As already noted, related issues are addressed by Franses and Paap (1999), but they concentrate primarily on the transition probabilities in the Markov switching model rather than the regime inferences of interest to us.

Our analysis uses the linear filters corresponding to the quarterly version of X-11, as most applications of the Markov switching model to the business cycle have been conducted at this frequency. We consider both the two-sided linear approximation to the X-11 procedure, and also the one-sided filter that is more appropriate to the real-time identification of the current state of the business cycle.

3.1 Data generation process

The data generating process (DGP) used in our experiments is the two-state regime-dependent model :

$$y_t = \mu_0 + \mu_d s_t + \varepsilon_t \quad (11)$$

where $\varepsilon_t \sim NID(0, \sigma^2)$ and s_t represents a binary state process capturing business cycle phases, with $\mu_d > 0$. For all simulations, estimation of (11) is undertaken treating the state as being unobserved, and making a Markov switching assumption for s_t . Thus, the researcher assumes that regime derives from a first-order Markov process with constant transition probabilities

$$\begin{aligned} P(s_t = 1 | s_{t-1} = 1) &= p & P(s_t = 0 | s_{t-1} = 1) &= 1 - p \\ P(s_t = 0 | s_{t-1} = 0) &= q & P(s_t = 1 | s_{t-1} = 0) &= 1 - q \end{aligned} \quad (12)$$

Obviously, since the DGP in (11) has no dynamics, no autoregressive lag augmentation is required when the Markov switching model is estimated using these data. However, this argument only applies when estimation is based on the unfiltered series y_t , since seasonal adjustment induces serial correlation, as discussed in Section 2. In order to control for any autocorrelation, we assume that a researcher working with quarterly data (seasonally adjusted or unadjusted) uses a four-lag autoregressive augmentation³. Thus, our Monte Carlo experiments are based on estimating the model:

$$\phi(L)y_t = \delta_0 + \delta_d s_t + e_t \quad (13)$$

with $\phi(L)$ being a fourth-order polynomial in the lag operator and e_t is a disturbance term which the researcher assumes to be an *iid* normal variate. This specification implies that the switching process changes the intercept, rather than the mean, of y_t , and has been considered by, among many others, Hamilton (1990). The use of the intercept-switching model instead of the mean-switching model is more practical in the context of a large Monte Carlo study, because in the mean-switching model the conditional distribution of y_t to be evaluated in the estimation algorithm would depend on four lags of the state process.

As already noted, we undertake simulations both where the regime in the DGP of (11) is deterministic and where it is stochastic.

³ It is possible the results of our simulations are to an extent sensitive to the choice of lag length for $\phi(L)$. In a sense, equation (13) is misspecified in that $\phi(L)$ will only approximately account for the autocorrelation induced by the application of the seasonal adjustment filter. Nonetheless, a researcher will in practice face a similar situation when using seasonally adjusted data. Thus, on the grounds of practicality, we choose as benchmark the lag length used by Hamilton (1989) and others in their applications of MS models to the business cycle.

3.1.1 DGP based on deterministic regimes

For the case of deterministic regimes, we assume that quarterly output data are available for the period 1951:I to 1996:IV, yielding $T=183$ values for the growth rate y_t . In this set of simulations, the state variable s_t is derived from the National Bureau of Economic Research (NBER) business cycle chronology for the US, with $s_t = 1$ in expansions and $s_t = 0$ in recessions over the period 1951-1996⁴. There are five distinct NBER recession regimes in this period, with durations between two and five quarters. Therefore, based on the analysis of Section 2, we anticipate seasonal adjustment will have some undesirable consequences for regime identification.

In terms of the parameters of (11), we set $\mu_0 = -0.5$ and $\mu_d = 1.2$, implying that growth is expected to decline at an annual rate of approximately 2 percent in recessions, with growth at an annual rate of approximately 2.8 percent in expansions. The disturbance standard deviation is set at $\sigma = 1$.

3.1.2 DGP based on stochastic regimes

In order to broaden our analysis, and to verify that the results obtained using deterministic regimes are not specific to that case, we also use the Markov switching model, (11) with (12), as the true DGP. In this case, $T=160$ observations are generated per each replication. When the underlying regime is stochastic, we assume that it is characterised by the transition probabilities $p = 0.9$ and $q = 0.65$, which are fairly representative values for quarterly data.. However, it is plausible that the effects of seasonal adjustment on regime identification may vary depending on how distinct are the two underlying regimes. In the context of the model of (11), the essential distinction between business cycle recessions and expansions is the

⁴ Following the usual convention in the use of NBER dates, the new regime is assumed to start at the beginning of the quarter after the month of the turning point.

difference between their mean growth rates, μ_d , and we consider a range of values for this parameter from 1.2 to 2.5 (with $\sigma = 1$ for both regimes). Thus, the value $\mu_d = 1.2$ used in the deterministic regime simulations is the smallest value considered in the stochastic regimes case. In general, since larger values of μ_d correspond to cases where the two regimes are more distinct, we anticipate that regime identification will improve as μ_d increases. However, the key issue in our context is whether the relative performances of seasonally adjusted and unadjusted data change with μ_d .

3.2 Regime probabilities

3.2.1 Full sample smoother probabilities

As discussed above, our main aim is to assess the impact of seasonal adjustment on regime inference in the Markov switching model. One set of regime probabilities of interest are then given by $P(s_t = j | \mathbf{Y}_T; \hat{\theta})$, which are the *full-sample smoother probabilities* that deliver the optimal probabilistic inference that state j applied at time t , based on complete sample information on y_t to time T (Hamilton, 1989). These probabilities are a by-product of parameter estimation, with $\hat{\theta}$ being the vector of maximum likelihood parameter estimates for the Markov switching model.

Based on each set of observations ($t = 1, \dots, T$) generated from the DGPs defined in subsections 3.1.1 and 3.1.2., the two-sided seasonally adjusted series is obtained as:

$$y_t^F = v_2(L)y_t, \quad t = 30, \dots, T-29 \quad (14)$$

where $v_2(L)$ is the linear filter approximation of the quarterly X-11 program, analysed in Section 2. The initial and final observations are not used due to the observations lost at each end when applying the two-sided filter. The unfiltered series uses observations over the same time period, so that:

$$y_t^{UF} = y_t, \quad t = 30, \dots, T-29 \quad (15)$$

In both cases, an additional four observations are used to create the autoregressive lags for the estimated model of (13), so that the smoothed probabilities relate to observations $t = 34, \dots, T-29$. We then compare the two sets of probabilities $P(s_t = j | Y_T^F; \hat{\theta}_F)$ and $P(s_t = j | Y_T^{UF}; \hat{\theta}_{UF})$, with (of course) separately estimated parameter vectors in the two cases. In the case of stochastic regimes, the performances of filtered and unfiltered data cannot be compared for specific regimes. Therefore, we adopt the widely-used summary measure of the quadratic probability score (QPS) relating to the regimes. When the smoother probabilities are used with filtered data, QPS is given by:

$$\text{QPS} = \frac{2}{(T_U - T_L + 1)} \sum_{t=T_L}^{T_U} [P(s_t = 1 | Y_T; \hat{\theta}_F) - s_t]^2 \quad (16)$$

with T_L and T_U being (respectively) the lower and upper sample observations for which regime inferences are obtained. A corresponding definition applies when unfiltered data are used.

3.2.2 Filter probabilities

The ex-post dating of business cycle regimes via the full-sample smoother is an important application of Markov switching models. Nonetheless, the use of such models in real time to assess the current state of the business cycle as new data become available, and to the forecast future regimes, is arguably of more practical interest. In addition, the properties of seasonal adjustment in a real time application are different from those discussed in the previous subsection, which assumes that enough data are available after the period of interest for the application of a symmetric adjustment filter.

The real-time tracking of the business cycle requires inference on the current regime employing only current and past information and can be accomplished also by Hamilton's

non-linear filtering algorithm, using the *filter probabilities* $P(s_t = j | Y_t; \hat{\theta})$, where the information set $Y_t = \{y_1, \dots, y_t\}$ includes only current and past information.

In addition, when considering the real time issue of regime identification using the filter probabilities, we must ensure that no future observations are employed in seasonal adjustment. Thus, we cannot use a two-sided seasonal adjustment filter throughout the sample period, as some observations in Y_t^F would then partly depend on information available only after period t . As already noted, the most recent seasonally adjusted observation is obtained from a one-sided filter, with preceding observations using asymmetric filters until $t - m$, with the full two-sided filter applied to earlier observations. However, for simplicity within our Monte Carlo simulation, all data points relevant for the computation of the filter probability are adjusted using the one-sided X-11 filter. Although this does not fully represent seasonal adjustment procedures used in practice, it does provide us with a guide as to how seasonal adjustment affects the recognition of a regime change in the period when the change occurs.

The filter probabilities $P(s_t = j | Y_t; \hat{\theta})$ require that parameter estimates be available. For this purpose, we assume the use of a historical dataset with parameter estimates obtained using observations $t = 30, \dots, T-29$ and with adjusted data for this estimation period derived from application of the two-sided linear filter $v_2(L)$, as in subsection 3.2.1. To capture the real-time aspect in the case of the DGP with deterministic regimes, we extend the sample period for which the DGP of (11) is applied to $t = 1, \dots, 310$, with the true state for the additional observations given by the NBER chronology repeated after 1996:IV, so that $s_{t+183} = s_t$ ($t = 1, \dots, 127$). Real-time regime inference is simulated by computing, without any further parameter estimation, the filter probabilities $P(s_t = j | Y_t; \hat{\theta})$ for 100 observations⁵ of the

⁵ Due to practical issues in interfacing the software programs used for the simulations, this analysis is restricted to 100 observations.

additional period, with filtered data for this period obtained using the one-sided filter $v_l(L)$. Allowing for autoregressive lags, the filter probabilities are computed for both adjusted and unadjusted data corresponding to $t = 215, \dots, 310$, with this period selected so that (with repetition of the NBER chronology) the five distinct historical recession regimes are included.

When examining filter probabilities in the case of the DGP with stochastic regimes, the same issues arise as in the context of the NBER dates. Our solution is the same as in that case, namely to estimate the parameters of the model based on one set of data (after application of the symmetric two-sided X-11 filter in the seasonally adjusted case) and to use these estimates to compute filter probabilities for the subsequent dataset, after application of the one-sided X-11 filter in the seasonally adjusted case. The sample observations used here are the same as outlined in the case of the NBER dates. Since the regimes used here are stochastic, the performances of filtered and unfiltered data are measured through the quadratic probability score (QPS) in an analogous way to that described in the previous subsection for the full sample smoother.

4. Simulation Results

All experiments employ 10,000 replications. Issues relating to estimation (including computer simulation) are discussed in the Appendix. We discuss first regime identification results related to the use of the two-sided symmetric X-11 filter, followed by those for the one-sided filter.

4.1 Symmetric Seasonal Adjustment Filter

4.1.1 DGP based on deterministic regimes

Beginning with results based on the NBER business cycle chronology, Figure 4 summarises the empirical distributions for the full-sample smoother probabilities for the filtered and

unfiltered data sets corresponding to the recession (lower) regimes, namely $P^F = P(s_t = 0 | \mathbf{Y}_T^F; \hat{\theta}_F)$ and $P^{UF} = P(s_t = 0 | \mathbf{Y}_T^{UF}; \hat{\theta}_{UF})$, with vertical lines denoting the true lower regime observations corresponding to NBER recessions. The distributions of estimates are summarized by graphing the third quartile, median and first quartile values corresponding to each quarter.

The third quartile, shown in panel (a), indicates that probabilities obtained using filtered data tend to point to the onset of recession too early and delay the recognition of recovery, with these features being more marked than with the use of unfiltered data. In other words, the filtered data obscures the dates of the regime changes, as predicted by the analysis of Section 2. However, while our analysis above indicates a symmetry for the distortion prior to the beginning of the regime and subsequent to its completion, the comparison of the estimated regimes for the Markov switching models using filtered and unfiltered data are asymmetric in Figure 4(a), with filtered data capturing the start of the regime more adequately than its end. This asymmetry presumably results from an interaction of the effect of the seasonal adjustment filter and the regime inference within the Markov switching model.

At the same time, however, it is also clear that once the upper (expansion) regime is well under way, the third quartile lower regime P^F probabilities are generally (and correctly) closer to zero than their unfiltered counterparts. Therefore, filtering reduces the chances that low-valued observations from the upper regime are mistakenly attributed to the lower regime. Within the lower regime, the filtered estimates are, on average, a little closer to the true regime value of one.

Turning to the plots of the median and first quartile of the smoother probabilities, panels (b) and (c) respectively of Figure 4, we concentrate first on the errors occurring during the lower regime. Here, filtering has the desirable effect of signalling periods in the lower regime more strongly than the unfiltered data probabilities. This effect is most noticeable

when the lower regime lasts a year or longer (the lower regimes commencing after $t = 74$, $t = 90$ and $t = 122$ in the graph). As illustrated in Figure 3, our analysis of Section 2 predicts that seasonal adjustment will have a negative effect on the detection of a short-lived regime of duration less than a year. However, when combined with the effects of the full-sample smoother probabilities, the median and first quartile probabilities during the short recession of two quarters (commencing in $t = 114$) are very similar whether filtered or unfiltered data are used. Further, the Markov switching model may completely miss genuine recessions. More specifically, when the recession duration is a year or less, the first quartile values in Figure 4(c) never rise above 0.5 during the recession whether adjusted or unadjusted data are employed, and they do not reach 0.5 using unadjusted data even when the recession duration is five quarters. Indeed, within the short recession of two quarters, the median probability barely reaches 0.5 for either type of data.

To sum up, filtering tends to make the full-sample smoother probabilities more inert and hence smoother. This is entirely consistent with the conclusion of Franses and Paap (1999) that the symmetric seasonal adjustment filter results in estimated transition probabilities, \hat{p} and \hat{q} , significantly larger than the true values and larger than their unfiltered counterparts. Within a regime (especially those of longer duration), this property of filtering facilitates correct regime identification, but it nevertheless acts against correct identification of regime switches. Thus, away from turning points, there is less scope for “false alarms” of a regime switch in the smoother probabilities with filtered data, so that the errors in P^F are smaller relative to P^{UF} once the upper regime is under way. However, the cost is particularly felt in the belated recognition of the beginning of a recovery from recession. Indeed, in the case of the two NBER recessions in the early 1980s (occurring between $t = 114$ and $t = 130$ in the figure), the third quartile probability for the filtered data never recognises the intermediate

recovery and implies continuous recession for approximately four years, from the end of 1979 until late 1983.

4.1.2 DGP based on stochastic regimes

In order to characterise these findings in a broader context, and also to examine the implications for measures of regime tracking, we turn to the Markov switching DGP and compute the QPS of (17) for each replication to quantify the comparison of the full-sample smoother probabilities using filtered and unfiltered data. In addition to the overall QPS, we compute separate values for observations in the upper regime and lower regimes, classified according to the actual s_t .

Figure 5 shows selected quantiles of the empirical distribution of QPS across the different values selected for μ_d in the DGP. It is clear from the overall QPS in panel (a) that the tracking of regimes is unambiguously worse with filtered data. Notice how the underperformance of filtered data is roughly constant (in absolute terms) across the different values for μ_d when considering the median QPS (denoted Q50), but this is not so for the upper quantiles of the distribution. Therefore, the potential for large errors in detecting the true underlying regime grows for filtered data in comparison with unfiltered data as the means in the two regimes become closer.

In the QPS computed for observations in the upper regime only, shown in panel (b), the filtered data fares even worse than in the overall QPS. Thus, the longer persistence of high values of P^F after the recession regime has come to an end, as noted above, plays a more important role in the final outcome than the better fit of P^F once the expansion regime is well under way. The deterioration of the fit of P^F relative to P^{UF} again substantially increases as the means in the two regimes become closer, indicating that the identification of the upper

regime becomes more problematic with filtered data (compared with unfiltered data) as the regimes become more similar.

In contrast, during the lower (recession) regime the fit obtained using filtered data is less markedly worse relative to unfiltered counterparts for large values of μ_d , with both probabilities producing roughly the same QPS values (in median terms) when $\mu_d = 1.8$. For a smaller difference between the means in the two regimes, the fit during the lower regime using filtered data is better than with unfiltered data, which is compatible with the pattern in Figure 4 for simulations using the NBER reference dates. Also, although the relative advantage of the P^F probabilities tends to grow with smaller values of μ_d , at the higher 95 and 90 percent quartiles, this advantage levels off, unlike the advantage for unfiltered data for the overall QPS and the QPS in expansions. It should also be noted that the QPS quantile values during recessions are always substantially greater than those during expansions (note the different scales used for the vertical axes in the three panels of Figure 5), again pointing to the possibility that the Markov switching model may miss genuine recessions.

The top panel of Table 1 presents a numerical comparison of the median values for the QPS from the full-sample smoother using adjusted and unadjusted data. It is again clear that the effects of filtering are not negligible. The median overall QPS is up to 21 percent larger with adjusted data, while the upper regime median QPS can be more than 50 percent higher (when $\mu_d = 1.2$). The relative effect is less marked in the lower regime, varying from a 17 percent deterioration to a similar improvement for adjusted compared with unadjusted data, depending on μ_d .

4.2 One-Sided Seasonal Adjustment Filter

4.2.1 DGP based on deterministic regimes

Using data based on the NBER business cycle regimes, the results for various quantiles of the empirical distribution of the recession filter probabilities $P(s_t = 0 | \mathbf{Y}_t^F; \hat{\theta}_F)$, after the application of the one-sided linear X-11 filter, and the corresponding probabilities using unadjusted data, $P(s_t = 0 | \mathbf{Y}_t^{UF}; \hat{\theta}_{UF})$, are presented in Figure 6.

It is clear that the filter probabilities and the use of the one-sided preliminary estimates of the seasonally adjusted data amplify the delay in recognising the onset of the upper (expansion) regime, compared with the effect noted in our full-sample smoother analysis. Again, this is particularly noticeable if one concentrates attention on the third quartile of the distribution; implying that there is a 25 percent chance of obtaining distortions like those seen in panel (a) of Figure 6. Indeed, this third quartile of the filter probabilities based on adjusted data does not yield a stable low chance of being in the lower regime (and hence a high probability of being in the upper regime) until ten to fifteen quarters into the upper regime. Also, it is very clear how the seasonal adjustment filter produces seasonal spikes in the filtered probabilities after the lower regime comes to an end; for instance, see its behaviour around $t = 231$. Notice, however, that the adjusted data probabilities do not signal the beginning of the lower regime earlier than the unfiltered data, as occurred at this quartile in the full sample smoother case. It appears, therefore, that effect was due to the two-sided nature of the adjustment undergone by the data fed into the smoother, rather than any direct effect of using adjusted data on the estimation of the parameters of the Markov switching model.

In our full-sample smoother analysis, the P^F probabilities during the lower regime were shown to be signalling the occurrence of that regime more clearly than those computed with unfiltered data. With the filter probabilities, however, this effect is much less noticeable

in panels (b) and (c) of Figure 6. With or without seasonal adjustment, the first quartile values in panel (c) of this figure point to the strong chance that the occurrence of a recession will be missed by the Markov switching model in real time (using filter probabilities).

The filter probabilities are also the basis of regime forecasting in the MS model. More precisely, the one-step-ahead prediction for the unobserved regime process can be computed as:

$$P(s_{t+1} = 0 | Y_t^F; \hat{\theta}_F) = P(s_t = 0 | Y_t^F; \hat{\theta}_F) \hat{q}_F + P(s_t = 1 | Y_t^F; \hat{\theta}_F) (1 - \hat{p}_F) \quad (17)$$

for adjusted data, with an analogous expression for unadjusted data using Y_t^{UF} and $\hat{\theta}_{UF}$. The empirical distribution of these one-step-ahead probabilities is summarised in Figure 7 for the NBER business cycle dates. Because the outcome of (17) is driven mainly by the filter probabilities and there is an upward bias in \hat{q}_F relative to \hat{q}_{UF} (Franses and Paap, 1999), the main conclusions are similar to those obtained for the filter probabilities. Once again, an important effect is that the use of adjusted data yields a higher probability for the continuation of the lower regime. On the one hand, this regime forecast yields greater errors when the switch to the upper regime occurs, but the errors will be smaller when the lower regime continues. From the analysis of Dacco and Satchell (1999), we might anticipate that the former effect will outweigh the latter when forecasting using the Markov switching model, since they show that incorrect regime forecasts are an underlying cause of the poor forecasting performance of regime-switching models.

4.2.2 DGP based on stochastic regimes

To examine these effects in a more general context, we again use QPS to quantify the impact of one-sided seasonal adjustment on regime detection in terms of the QPS measure for the case where the DGP is a Markov switching model with μ_d ranging from 1.2 to 2.5. Here the

QPS is computed for both the filter and the one-step-ahead regime probabilities. To assess the effects on regime forecasting, we compute QPS using (16), but where the regime refers to $t+1$ with the prediction based on the information set to t . The QPS measures for each regime separately, according to the outcome s_{t+1} in each replication, are also computed. Quantiles of QPS for the estimated filter probabilities and one-step ahead prediction probabilities are shown in Figures 8 and 9 respectively.

In overall QPS terms, using adjusted data corrupts the regime tracking ability of the filter probabilities substantially more than was the case in the full-sample smoother analysis, especially in terms of increasing the chances of relatively large errors, with this being more noticeable as the regime means become closer. This is due, as in the full-sample smoother case, to a substantially worse fit during the expansion regime, which stems mainly from a delay in the recognition of the regime's onset. As seen above, this delay is much more accentuated in the case of the filter probabilities as compared to the outcome of the full-sample smoother. On the other hand, during the lower regime the fit using adjusted data is almost always better than using unadjusted data. Still, this does not offset the bad fit during the upper regime, hence the overall worse results obtained using adjusted data.

The conclusions from the one-step-ahead regime probabilities (Figure 8) are generally similar. The main difference is that within the lower regime, the fit using filtered data is superior to that using unadjusted data for all quantiles and μ_d values shown. This is presumably due to the extra role of the lower regime transition probability estimate in (17), where the positive bias in \hat{q}_F further accentuates the effects of using adjusted data for the better fit obtained during the lower regime. However, the overall QPS is still always worse for filtered data for the cases considered.

The second and third panels of Table 1 present the median QPS values underlying Figures 8 and 9, for the filter and one-step ahead regime probabilities respectively. The filter

probabilities using adjusted data show a deterioration of up to 33 percent, with the corresponding values being substantially larger when the upper regime QPS is considered. Perhaps surprisingly, the relative deterioration is smaller for the one-step-ahead probabilities, ranging between 5 and 7 percent for the overall QPS, but (as also seen in Figure 9) this is due to offsetting effects across the two regimes.

5. Concluding Remarks

This paper has considered the effects of seasonal adjustment on the detection of business cycle regimes, both analytically and when these regimes are identified through the use of a Markov switching model. The overall conclusion is that seasonal adjustment can distort the information about the extent and timing of the breaks (or turning points) that underlie regime identification.

Nevertheless, the picture is not entirely straightforward, because in some instances seasonal adjustment can have the effect of clarifying the regime. Indeed, through the smoothing inherent in seasonal adjustment, the filtered data tend to produce less false turning point signals, albeit they will detect the occurrence of actual turning points with more difficulty. As measured by summary statistics of regime tracking like the Quadratic Probability Score, however, the filtering procedure deteriorates the fit of regimes in the Markov switching model overall, with this result being mainly dominated by a belated signal of the occurrence of a business cycle trough. Our analytical results shed light on why this occurs, with the effect of the X-11 seasonal adjustment filter being to reduce the magnitude of the breaks that mark regime changes. This reduction is larger when the break is later reversed within a relatively short period of time, as occurs when a recession lasts between (say) six and fifteen months, and it is realistic to anticipate that seasonal adjustment could reduce the magnitude of the regime breaks by the order of 20 percent.

Our Monte Carlo investigation shows that the effects on regime tracking are exacerbated when using the one-sided filters in X-11, such as those employed by statistical agencies to produce preliminary releases of seasonally adjusted series. Again, the most prominent effect is a delay in signalling the beginning of the upper (expansion) regime. This has important implications for real-time regime identification purposes, as well as for regime forecasting. Because the use of seasonally adjusted data corrupts regime tracking specifically around turning points, this further suggests that researchers should be aware of the effects we have unearthed when using seasonally adjusted data for the detection of regime changes. Although illustrated here primarily in the context of Markov switching models, the analytical results imply that effects of this type for seasonal adjustment should be anticipated whatever method is used to identify business cycle turning points.

APPENDIX

Software and Procedures

In the simulations for the Markov switching model as the DGP, the initial state (s_0) is stochastic, but selected according to the unconditional state probability implied by the DGP, as

$$s_0 = \begin{cases} 0 & \text{when } x > \frac{1 - q}{2 - p - q} \\ 1 & \text{otherwise} \end{cases}$$

with x being randomly drawn for each replication from the (0, 1) uniform distribution. In estimation, the initial state probability is treated as known and equal to $(1 - q_0)/(2 - p_0 - q_0)$, where p_0 and q_0 are the transition probability values used to initialise the estimation of these parameters. In all cases, estimation is initialised with $p_0 = 0.9$ and $q_0 = 0.65$; for the Markov switching DGPs, these are the true parameter values.

All estimations of the Markov switching model are conducted using the GAUSS procedures available from the Bank of Canada (van Norden and Vigfusson, 1996). These have the advantage of allowing the use of analytical gradients in the Maximum Likelihood estimation of the Markov switching model. Estimation was undertaken by direct optimisation of the likelihood function; the EM procedure included in the Bank of Canada programs was not employed due to some problems in convergence. Optimisation is undertaken with the Broyden, Fletcher, Goldfarb and Shanno algorithm in the Gauss procedure MAXLIK (1995 version), with line search set to STEPBT. The convergence criterion is set to be a tolerance of 10^{-5} for the gradient of the estimated coefficients. All random draws were generated by the routines RNDU and RNDN contained in GAUSS 3.2.

An important practical difficulty in the estimation of Markov switching models is the presence of local optima in the likelihood function. Therefore, depending on the particular set up, the global maximum may be undetected in a substantial proportion of the Monte Carlo replications unless a large number of starting values are used to initiate the nonlinear maximisation routine in each replication. However, the use of a large grid of starting values severely limits the possibility of designing a large Monte Carlo study, due to the long computation times involved.

Hamilton (1991, 1996) suggests using the DGP parameter values as starting values for the maximisation iterations to avoid this problem. This approach works fairly well for the parameterisations of (11) that involve large differences between the means in the two regimes, but not so well for those parameterisations that imply a less dramatic (and perhaps more realistic) difference. Indeed, as the means in the two regimes became closer, the density of $\hat{\mu}_d$ became bimodal, with a global maximum at around μ_d and a local maximum at zero. This corresponds to a local maximum known to occur in the likelihood function for the Markov switching model (see, for instance, Hamilton, 1996). Therefore, starting our iterations at the true parameter values may produce seriously misleading results, as a large proportion of the replications get stuck in the local maxima around $\hat{\mu}_d = 0$, at least for those parameterisations that involve a small difference in means⁶.

A possible solution would be to run a grid of different starting values for those replications that deliver an estimate of $\hat{\mu}_d$ close to zero, but this would leave open the

⁶ Hansen (1992) finds a similar phenomenon in studying the finite sample behaviour of the likelihood ratio test for Markov switching nonlinearity. He warns against a very high chance that the estimation will end up at a local optimum if iterations are started at the true DGP parameter values under the null of a linear DGP.

question of how close to zero $\hat{\mu}_d$ should be to switch to this extensive grid search. Further experimentation, based on Garcia (1998), showed that most of our suspected local optima produced near zero likelihood ratios test values in comparison with a linear autoregressive model, with less than 0.01 being a convenient value. For these replications, a grid of 10 different starting values succeeded in moving $\hat{\mu}_d$ and the likelihood ratio away from zero in all cases⁷. On the other hand, our experiments showed that for replications with likelihood ratios above 0.01, no extensive grid search could find a larger value for the likelihood function than the one obtained from iterations started at the true parameter values. Thus all the reported results are conducted based on this starting value criterion⁸.

Seasonal adjustment is undertaken using the SAS/ETS for Windows 6.1 version of the additive X-11 algorithm with standard options, with the exception of outlier corrections, which is disabled. In the case of the one-sided filter, each element y_t^F , $t = 211, \dots, 310$, is obtained by recursive application of the X-11 program, that is, by running the SAS routine on a dataset containing observations only through to time t and fetching the last seasonally adjusted value. This procedure is repeated for 10,000 replications.

⁷ The grid used was $\delta_d = \mu_d^{DGP} + [0.2, 2]$ in steps of 0.2, with μ_d^{DGP} being the true value of this parameter in the DGP considered. Initial values of δ_0 , p and q were always at the DGP values. In some experiments, the grid of starting values was extended to these parameters, but this always produced the same estimates as the restricted grid search.

⁸ Note that for filtered data, the correspondence, $\delta_0 = \mu_0$, $\delta_d = \mu_d$ does not hold. Nevertheless, starting estimation at these values does not seem to affect the ability of the optimisation algorithm to find the global maximum.

REFERENCES

- Barsky, R.B., and Miron, J.A., 1989, “The Seasonal Cycle and the Business Cycle”, *Journal of Political Economy*, 97, 503-34.
- Bell, W.R., and Hillmer, S.C., 1984, “Issues Involved with the Seasonal Adjustment of Economic Time Series”, *Journal of Business and Economic Statistics*, 2, 291-320.
- Burridge, P., and Wallis, K.F., 1984, “Unobserved-Components Models for Seasonal Adjustment Filters”, *Journal of Business and Economic Statistics*, 2, 350-359.
- Canova, F., and Ghysels, E., 1994, “Changes in Seasonal Patterns - Are They Cyclical?”, *Journal of Economic Dynamics and Control*, 18, 1143-1171.
- Cecchetti, S.G., and Kashyap, A.K., 1996, “International Cycles”, *European Economic Review*, 40, 331-360.
- Christiano, L.J. and Todd, R.M. , 2002, “The Conventional Treatment of Seasonality in Business Cycle Analysis: Does It Create Distortions?”, *Journal of Monetary Economics*, 49, 335-364.
- Dacco, R., and Satchell, S., 1999, “Why Do Regime-Switching Models Forecast So Badly?”, *Journal of Forecasting*, 18, 1-16.
- Dagum, E.B., 1980, “The X-11-ARIMA Seasonal Adjustment Method”, Report E12-564E, Statistics Canada, Ottawa.
- Findley, D.F, Monsell, B.C., Bell, W.R., Otto, M.C., and Chen, B.-C., 1998, “New Capabilities and Methods of the X-12-ARIMA Seasonal-Adjustment Program”, *Journal of Business and Economic Statistics*, 16, 127-52.
- Franses, P.H. and Paap, R., 1999, “Does Seasonality Influence the Dating of Business Cycle Turning Points?” , *Journal of Macroeconomics*, 21, 79-92.

Garcia, R., 1998, “Asymptotic Null Distribution of the Likelihood Ratio Test in Markov Switching Models”, *International Economic Review*, 39, 763-88.

Ghysels, E., Granger, C.W.J. and Siklos, P. (1996), “Is Seasonal Adjustment a Linear or Nonlinear Data Filtering Process”, *Journal of Business and Economic Statistics*, 14, 374-386.

Ghysels, E. and Osborn, D.R., 2001, *The Econometric Analysis of Seasonal Time Series*, Cambridge, UK: Cambridge University Press.

Ghysels, E. and Perron, P., 1993, “The Effect of Seasonal Adjustment Filters on Tests for a Unit Root”, *Journal of Econometrics*, 55, 57-98.

Ghysels, E. and Perron, P., 1996, “The Effect of Linear Filters on Dynamic Time Series with Structural Change”, *Journal of Econometrics*, 70, 69-97.

Hamilton, J.D., 1989, “A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle”, *Econometrica*, 57, 357-84

Hamilton, J.D., 1990, “Analysis of Time Series Subject to Changes in Regime”, *Journal of Econometrics*, 45, 39-70.

Hamilton, J.D., 1991, “A Quasi-Bayesian Approach to Estimating Parameters for Mixtures of Normal Distributions”, *Journal of Business and Economic Statistics*, 9, 27-39

Hamilton, J.D., 1996, “Specification Testing in Markov-Switching Time-Series Models”, *Journal of Econometrics*, 70, 127-57

Hansen, B.E., 1992, “The Likelihood Ratio Test under Nonstandard Conditions: Testing the Markov Switching Model of GNP” *Journal of Applied Econometrics*, 7, S61-S82.

Krane, S. and Wascher, W., 1999, “The Cyclical Sensitivity of Seasonality in US Employment”, *Journal of Monetary Economics*, 44, 523-553.

Laroque, G., 1977, “Analyse d’une Méthode de Désaisonnalisation: Le Programme X-11 du US Bureau of the Census Version Trimstrielle”, *Annales de l’INSEE*, 88, 105-127.

Matas-Mir, A., and Osborn, D.R., 2003, “The Extent of Seasonal/Business Cycle Interactions in European Industrial Production”, Forthcoming in Monographs of Official Statistics, Statistical Office of the European Communities (Eurostat). Available at <http://europa.eu.int/comm/eurostat/Public/datashop/print-product/EN?catalogue=euroindicators&product=CO-osborn-EN&mode=download>

Matas-Mir, A., and Osborn, D.R., 2004, “Does Seasonality Change over the Business Cycle?: An Investigation using Monthly Industrial Product Series”, *European Economic Review*, Forthcoming.

Sims, C.A., 1974, “Seasonality in Regression”, *Journal of the American Statistical Association*, 69, 618-27.

van Norden, S. and Vigfusson, R., 1996, “Regime-Switching Models: A Guide to the Bank of Canada Gauss Procedures”, Working Paper 96-3, Bank of Canada.

Wallis, K. F., 1974, “Seasonal adjustment and relations between variables”, *Journal of the American Statistical Association*, 69, 18-31.

Table 1. Median Quadratic Probability Score for Simulations of Markov Switching DGP

Full sample smoother												
(a) Filtered data (X-11 symmetric filter)												
μ_d	Overall			(b) Unfiltered data			(c) Ratio (a)/(b)			Lower reg.		
	Upper reg.	Lower reg.	Overall	Upper reg.	Lower reg.	Overall	Upper reg.	Lower reg.	Overall	Upper reg.	Lower reg.	
2.5	0.128	0.065	0.268	0.106	0.051	0.228	1.211	1.271	1.211	1.271	1.176	
2	0.195	0.096	0.389	0.169	0.078	0.361	1.150	1.228	1.150	1.228	1.077	
1.8	0.228	0.114	0.442	0.204	0.090	0.440	1.116	1.275	1.116	1.275	1.003	
1.6	0.264	0.138	0.500	0.242	0.102	0.526	1.093	1.357	1.093	1.357	0.950	
1.4	0.307	0.170	0.556	0.285	0.118	0.622	1.078	1.444	1.078	1.444	0.893	
1.2	0.354	0.218	0.599	0.326	0.140	0.722	1.085	1.563	1.085	1.563	0.830	

Filter probabilities												
(a) Filtered data (X-11 one-sided filter)												
μ_d	Overall			(b) Unfiltered data			(c) Ratio (a)/(b)			Lower reg.		
	Upper reg.	Lower reg.	Overall	Upper reg.	Lower reg.	Overall	Upper reg.	Lower reg.	Overall	Upper reg.	Lower reg.	
2.5	0.153	0.095	0.295	0.115	0.059	0.247	1.330	1.593	1.330	1.593	1.194	
2	0.218	0.134	0.404	0.182	0.093	0.385	1.200	1.433	1.200	1.433	1.050	
1.8	0.248	0.150	0.450	0.213	0.107	0.450	1.164	1.408	1.164	1.408	1.000	
1.6	0.288	0.178	0.515	0.255	0.123	0.540	1.130	1.443	1.130	1.443	0.955	
1.4	0.326	0.203	0.568	0.294	0.138	0.624	1.110	1.469	1.110	1.469	0.910	
1.2	0.368	0.231	0.643	0.329	0.135	0.772	1.118	1.717	1.118	1.717	0.833	

One quarter ahead probabilities												
(a) Filtered data (X-11 one-sided filter)												
μ_d	Overall			(b) Unfiltered data			(c) Ratio (a)/(b)			Lower reg.		
	Upper reg.	Lower reg.	Overall	Upper reg.	Lower reg.	Overall	Upper reg.	Lower reg.	Overall	Upper reg.	Lower reg.	
2.5	0.288	0.126	0.824	0.273	0.109	0.835	1.054	1.154	1.054	1.154	0.986	
2	0.315	0.144	0.852	0.300	0.121	0.874	1.048	1.188	1.048	1.188	0.974	
1.8	0.330	0.150	0.852	0.314	0.125	0.892	1.051	1.199	1.051	1.199	0.956	
1.6	0.349	0.162	0.875	0.331	0.127	0.936	1.055	1.279	1.055	1.279	0.934	
1.4	0.368	0.177	0.873	0.347	0.134	0.954	1.060	1.322	1.060	1.322	0.915	
1.2	0.385	0.191	0.896	0.359	0.122	1.041	1.070	1.563	1.070	1.563	0.860	

Note: Median values are computed over 10,000 Monte Carlo replications

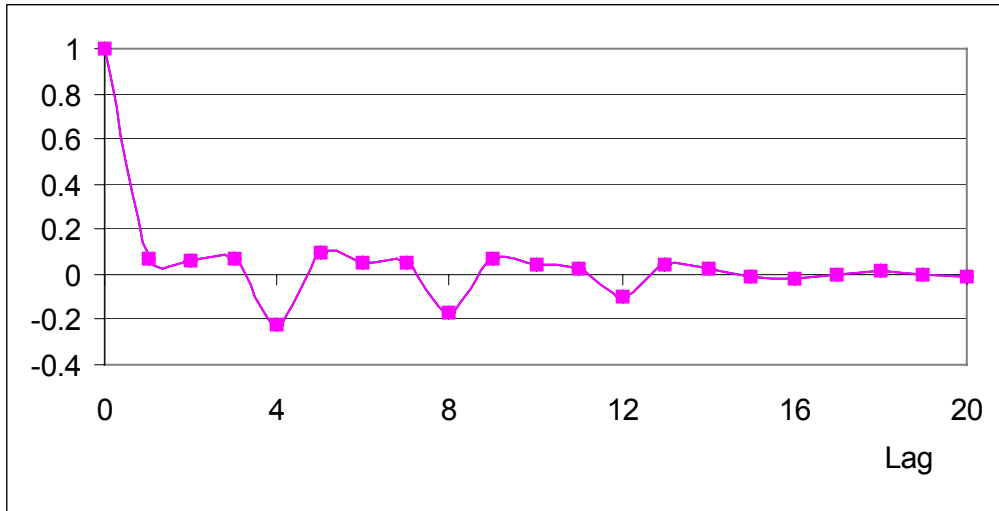
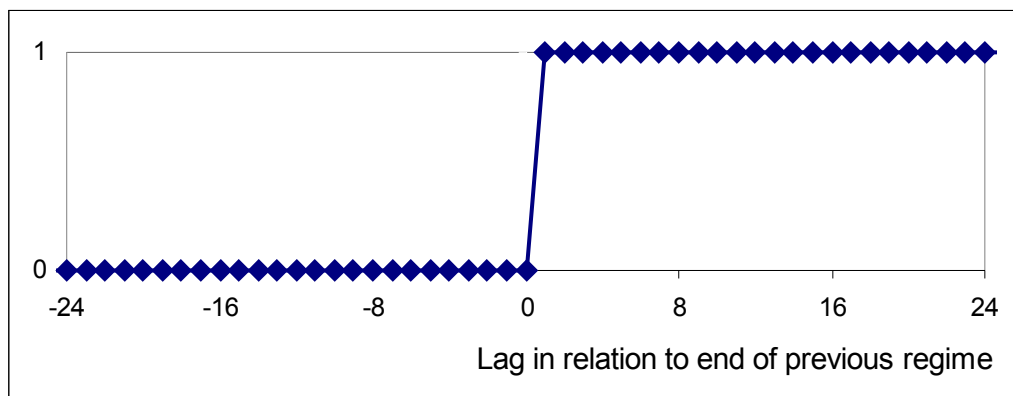


Figure 1. Autocorrelation function for white noise after seasonal adjustment by the two-sided linear X-11 filter

(a) Structural break series



(b) Distortion induced in break series by seasonal adjustment

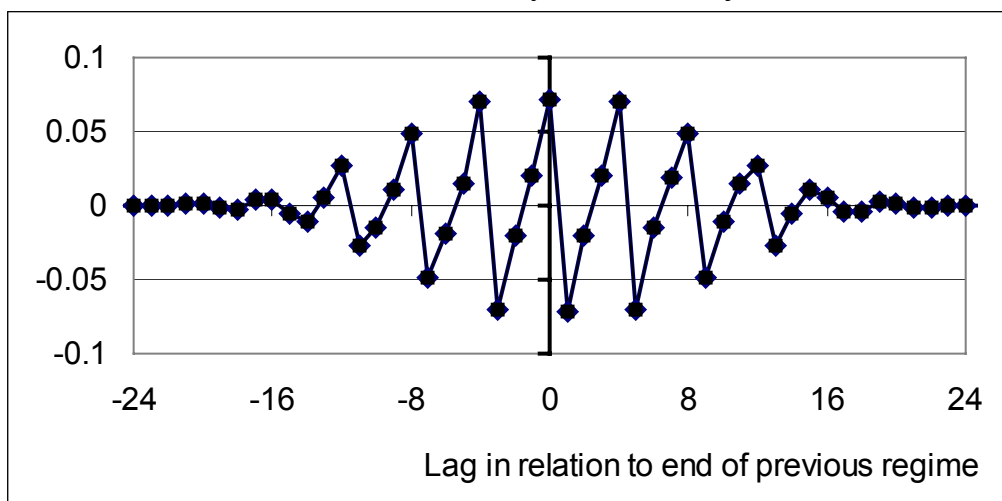
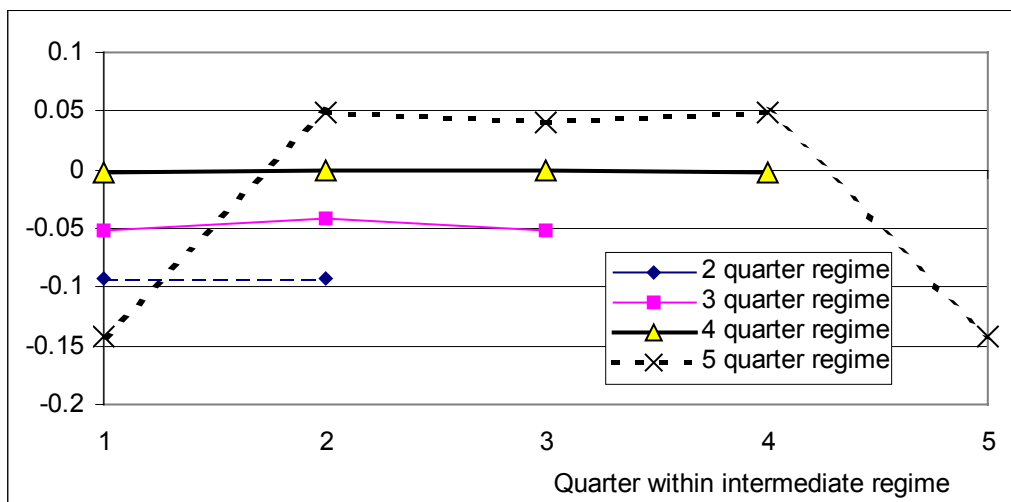


Figure 2. Distortion caused by applying the two-sided linear X-11 filter to a single structural break

(a) Distortion within intermediate regime



(b) Distortion from end of intermediate regime

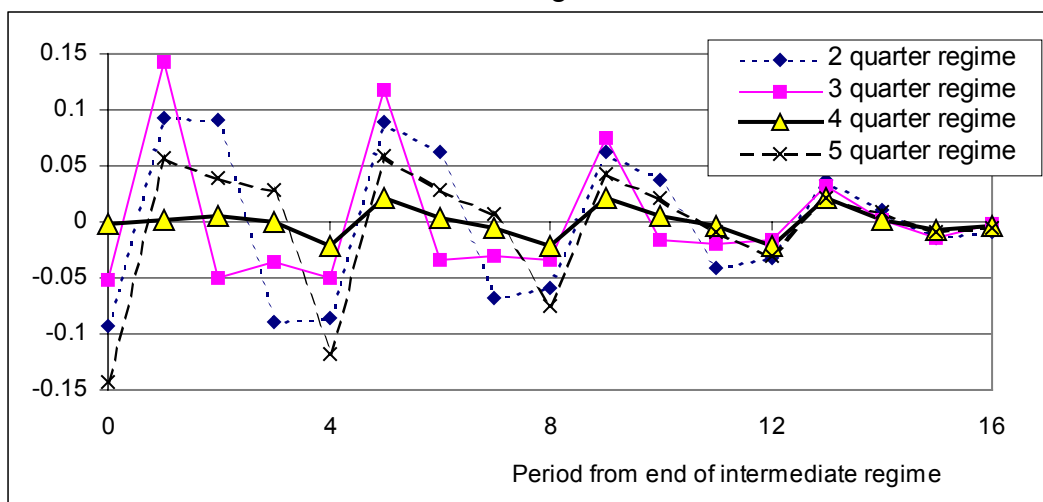


Figure 3. Distortion caused by applying the two-sided linear X-11 filter to a reversed structural break

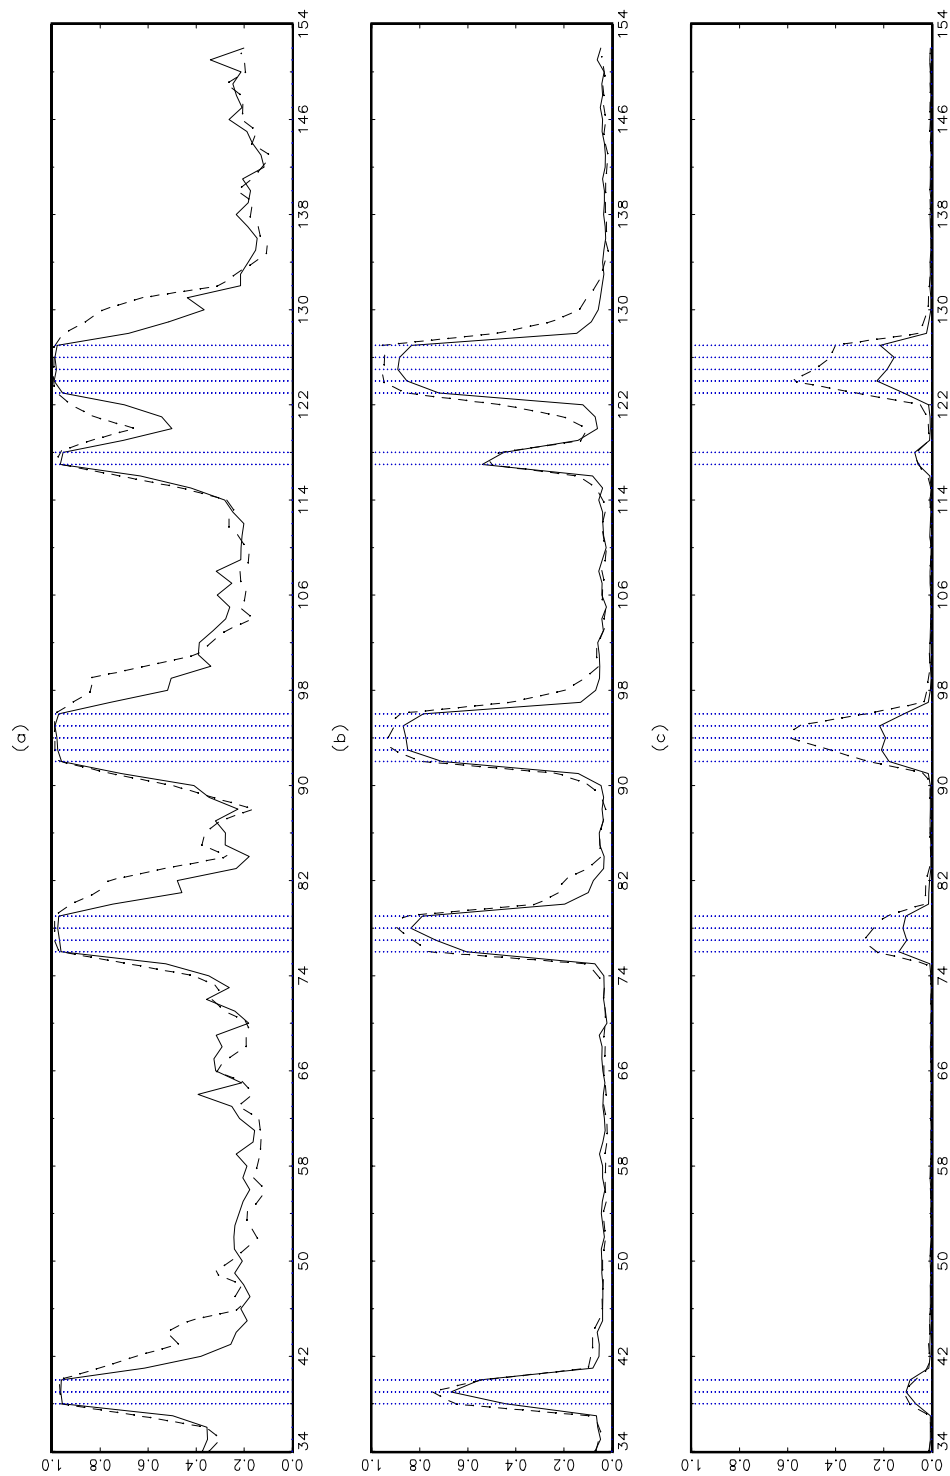


Figure 4. Markov-Switching full-sample smoother probabilities.

10000 replications. (a) Third Quartile; (b) Median; (c) First Quartile.

Solid line; raw data; dashed line : X-11 filtered data (symmetric filter) ; vertical lines : DGP's lower regime.

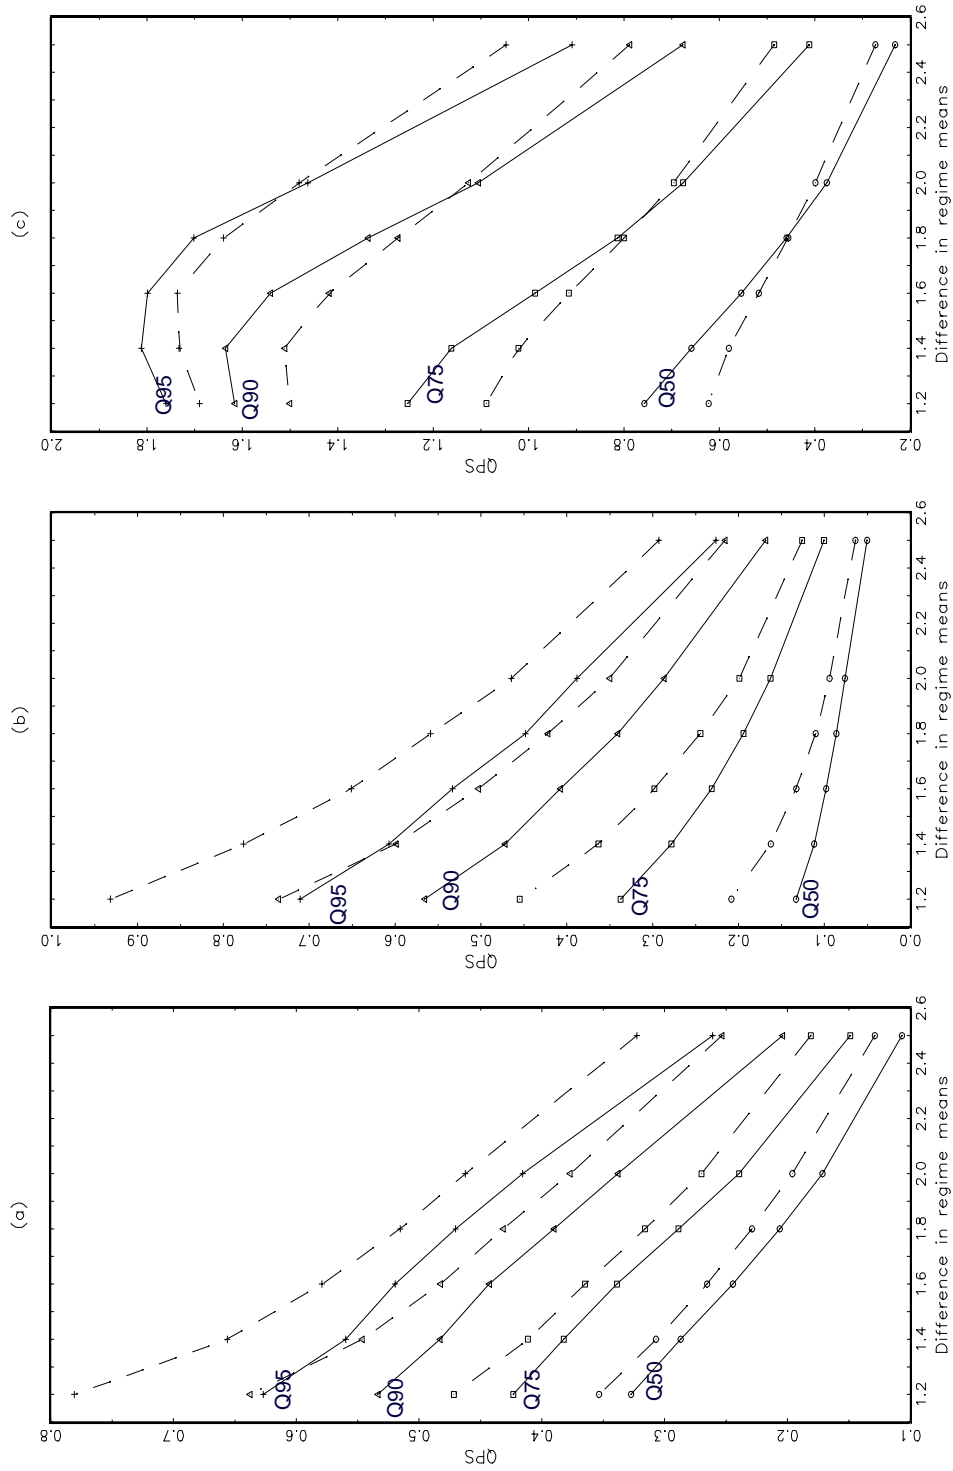


Figure 5. Markov-Switching full-sample smoother probabilities, QPS, selected quantiles. 10000 replications. (a) Overall QPS; (b) Upper regime QPS; (c) Lower regime QPS. Solid lines : unfiltered data; dashed lines: X-11 filtered data (symmetric filter).

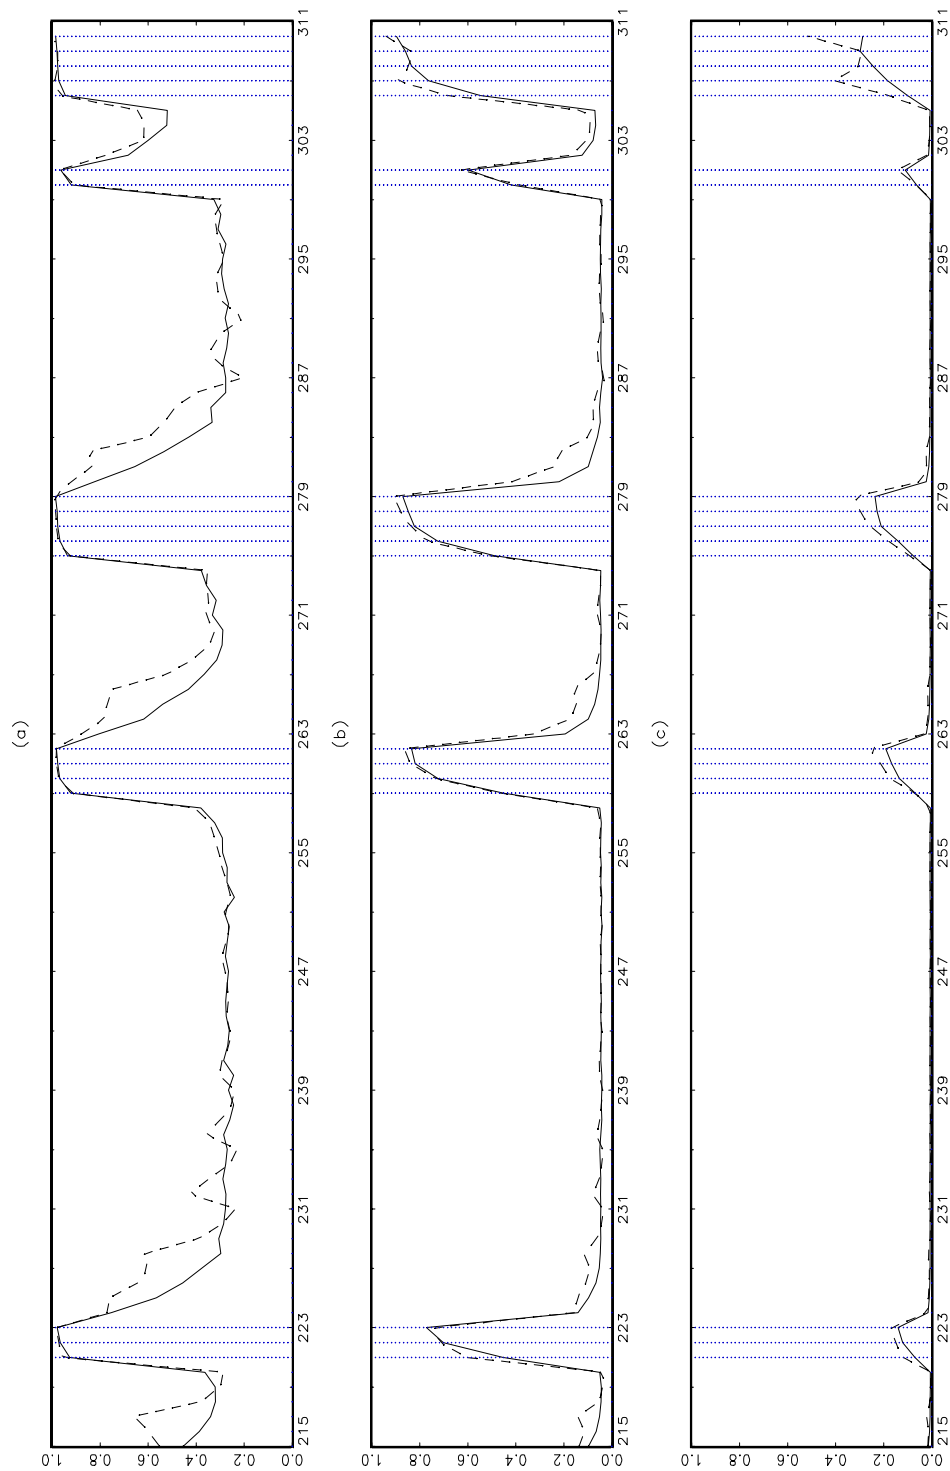


Figure 6. Markov-Switching filter probabilities.

10000 replications. (a) Third Quartile; (b) Median; (c) First Quartile.

Solid line : raw data; dashed line : X-11 filtered data (one-sided filter); vertical lines : DGP's lower regime.

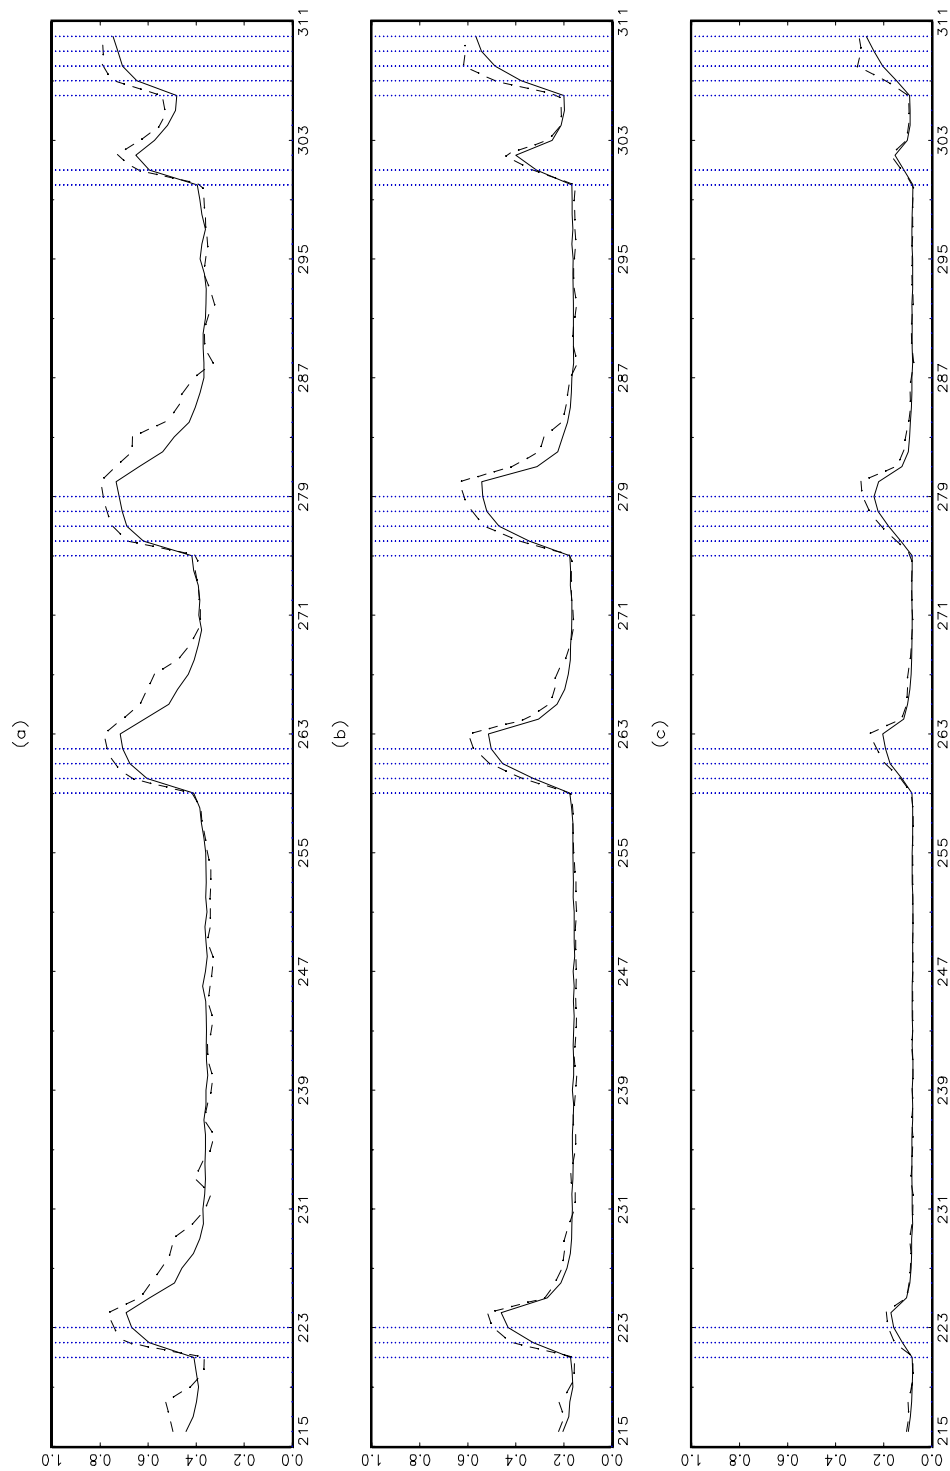


Figure 7. Markov-Switching one-step-ahead regime probabilities. 10000 replications. (a) Third Quartile; (b) Median; (c) First Quartile.

Solid line: raw data; dashed line : X-11 filtered data (one-sided filter); vertical lines : DGP's lower regime.

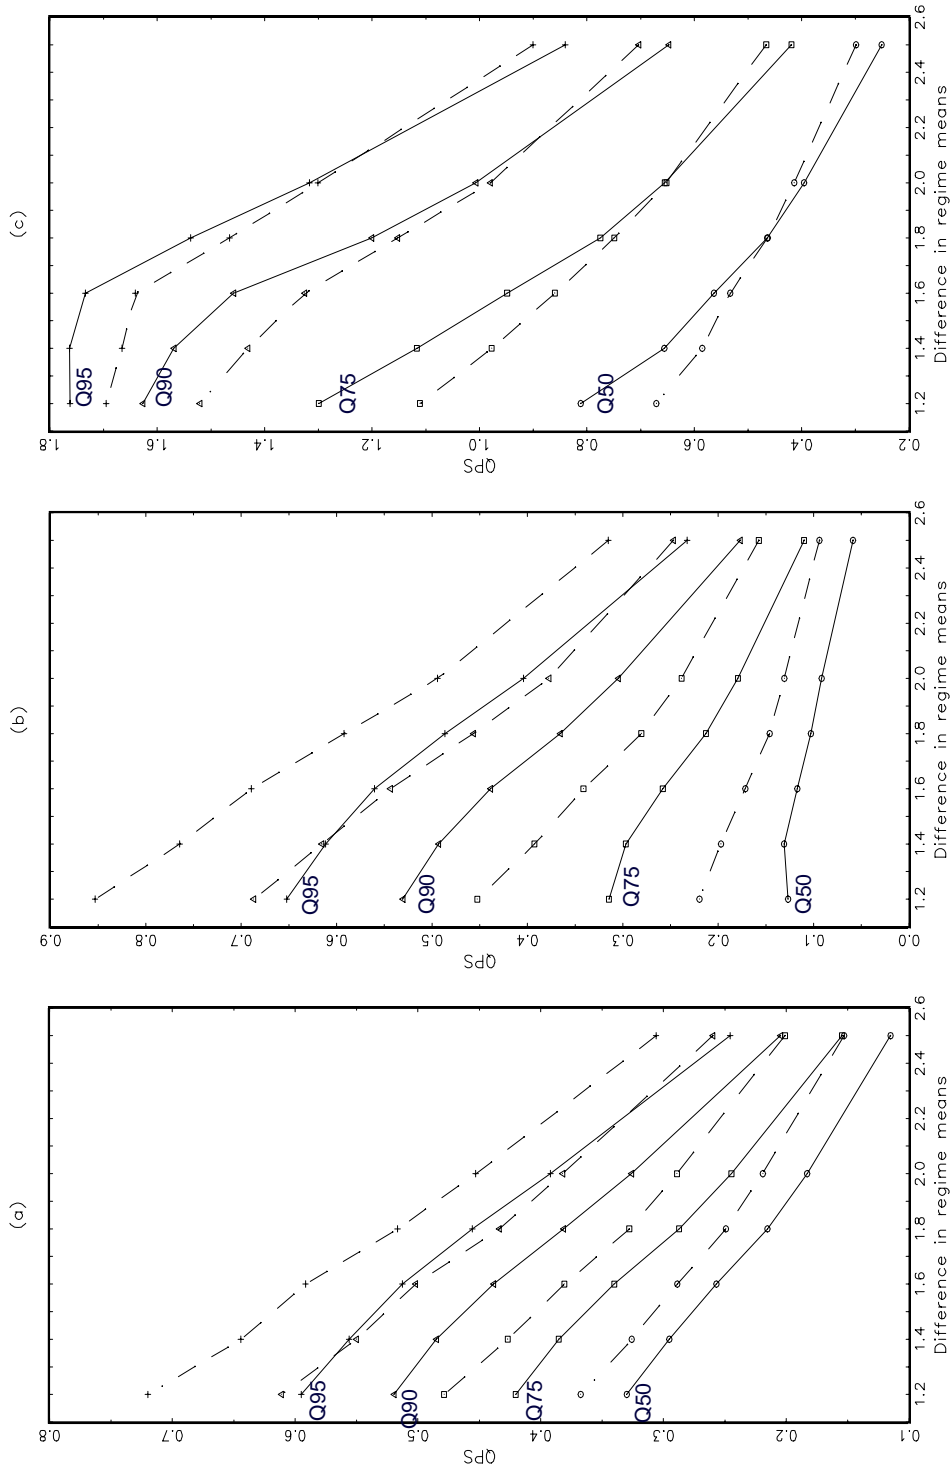


Figure 8. Markov-Switching filter probabilities, QPS, selected quantiles. 10000 replications. (a) Overall QPS; (b) Upper regime QPS; (c) Lower regime QPS. Solid lines : unfiltered data; dashed lines: X-11 filtered data (one-sided filter).

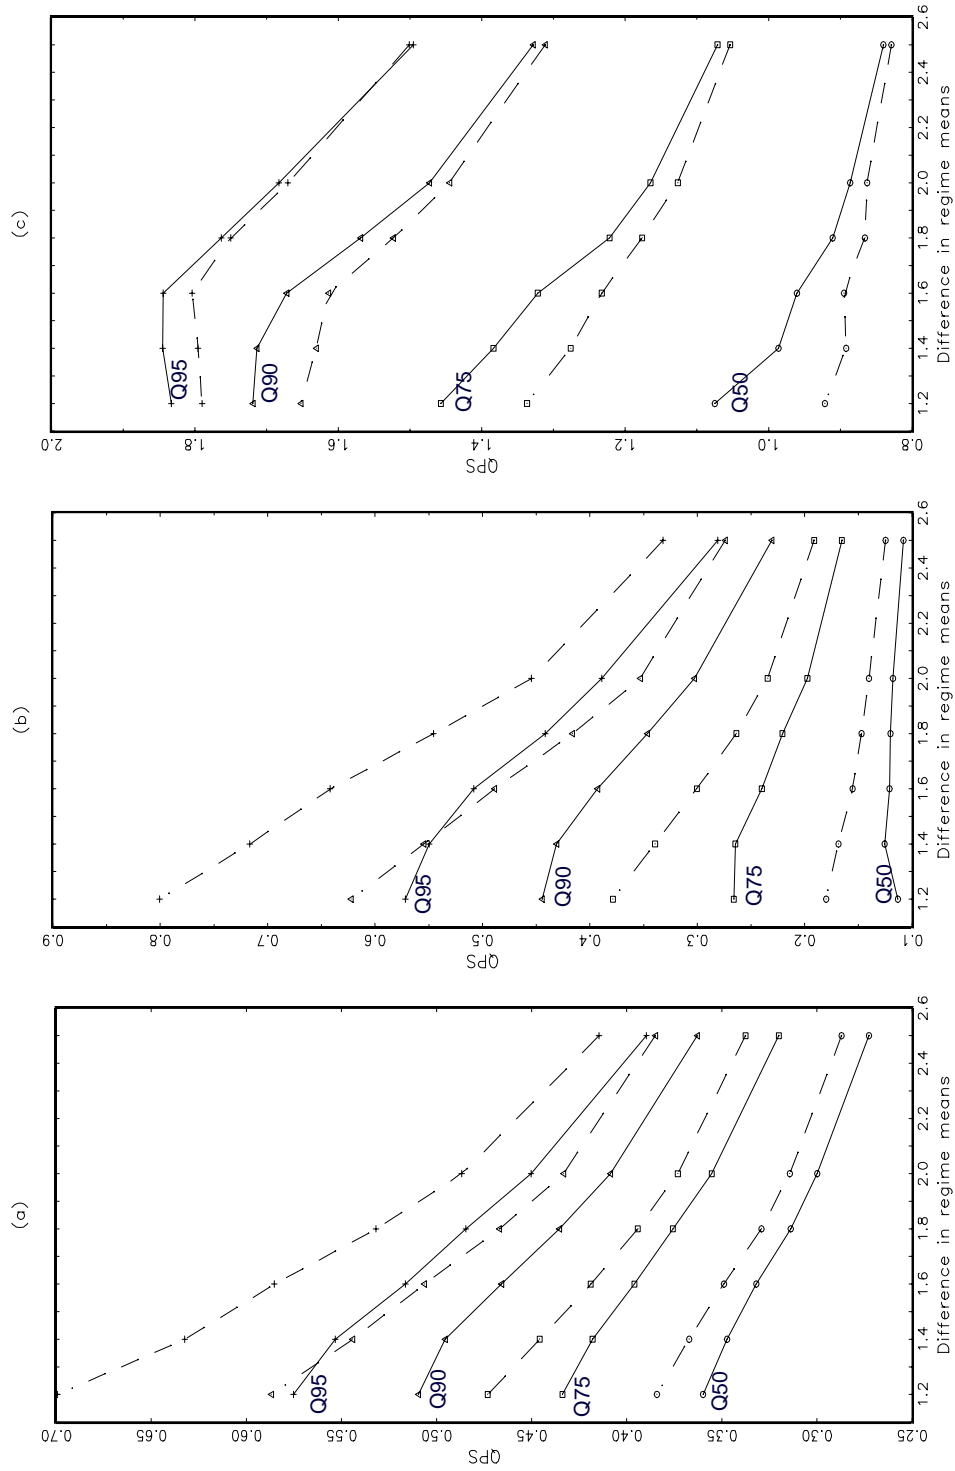


Figure 9. Markov-Switching one-step-ahead probabilities, QPS, selected quantiles. 10000 replications. (a) Overall QPS; (b) Upper regime QPS; (c) Lower regime QPS. Solid lines : unfiltered data; dashed lines: X-11 filtered data (one-sided filter).

European Central Bank working paper series

For a complete list of Working Papers published by the ECB, please visit the ECB's website (<http://www.ecb.int>).

- 302 "Deposit insurance, moral hazard and market monitoring" by R. Gropp and J. Vesala, February 2004.
- 303 "Fiscal policy events and interest rate swap spreads: evidence from the EU" by A. Afonso and R. Strauch, February 2004.
- 304 "Equilibrium unemployment, job flows and inflation dynamics" by A. Trigari, February 2004.
- 305 "A structural common factor approach to core inflation estimation and forecasting" by C. Morana, February 2004.
- 306 "A markup model of inflation for the euro area" by C. Bowdler and E. S. Jansen, February 2004.
- 307 "Budgetary forecasts in Europe - the track record of stability and convergence programmes" by R. Strauch, M. Hallerberg and J. von Hagen, February 2004.
- 308 "International risk-sharing and the transmission of productivity shocks" by G. Corsetti, L. Dedola and S. Leduc, February 2004.
- 309 "Monetary policy shocks in the euro area and global liquidity spillovers" by J. Sousa and A. Zaghini, February 2004.
- 310 "International equity flows and returns: A quantitative equilibrium approach" by R. Albuquerque, G. H. Bauer and M. Schneider, February 2004.
- 311 "Current account dynamics in OECD and EU acceding countries – an intertemporal approach" by M. Bussière, M. Fratzscher and G. Müller, February 2004.
- 312 "Similarities and convergence in G-7 cycles" by F. Canova, M. Ciccarelli and E. Ortega, February 2004.
- 313 "The high-yield segment of the corporate bond market: a diffusion modelling approach for the United States, the United Kingdom and the euro area" by G. de Bondt and D. Marqués, February 2004.
- 314 "Exchange rate risks and asset prices in a small open economy" by A. Derviz, March 2004.
- 315 "Option-implied asymmetries in bond market expectations around monetary policy actions of the ECB" by S. Vähämaa, March 2004.

- 316 “Cooperation in international banking supervision” by C. Holthausen and T. Rønde, March 2004.
- 317 “Fiscal policy and inflation volatility” by P. C. Rother, March 2004.
- 318 “Gross job flows and institutions in Europe” by R. Gómez-Salvador, J. Messina and G. Vallanti, March 2004.
- 319 “Risk sharing through financial markets with endogenous enforcement of trades” by T. V. Köppl, March 2004.
- 320 “Institutions and service employment: a panel study for OECD countries” by J. Messina, March 2004.
- 321 “Frequency domain principal components estimation of fractionally cointegrated processes”
by C. Morana, March 2004.
- 322 “Modelling inflation in the euro area” by E. S. Jansen, March 2004.
- 323 “On the indeterminacy of New-Keynesian economics” by A. Beyer and R. E. A. Farmer, March 2004.
- 324 “Fundamentals and joint currency crises” by P. Hartmann, S. Straetmans and C. G. de Vries, March 2004.
- 325 “What are the spill-overs from fiscal shocks in Europe? An empirical analysis” by M. Giuliadori
and R. Beetsma, March 2004.
- 326 “The great depression and the Friedman-Schwartz hypothesis” by L. Christiano, R. Motto and
M. Rostagno, March 2004.
- 327 “Diversification in euro area stock markets: country versus industry” by G. A. Moerman, April 2004.
- 328 “Non-fundamental exchange rate volatility and welfare” by R. Straub and I. Tchakarov, April 2004.
- 329 “On the determinants of euro area FDI to the United States: the knowledge-capital-Tobin's Q framework,
by R. A. De Santis, R. Anderton and A. Hijzen, April 2004.
- 330 “The demand for euro area currencies: past, present and future” by B. Fischer, P. Köhler and F. Seitz, April 2004.
- 331 “How frequently do prices change? evidence based on the micro data underlying the Belgian CPI” by
L. Aucremanne and E. Dhyne, April 2004.
- 332 “Stylised features of price setting behaviour in Portugal: 1992-2001” by M. Dias, D. Dias
and P. D. Neves, April 2004.
- 333 “The pricing behaviour of Italian firms: New survey evidence on price stickiness” by
S. Fabiani, A. Gattulli and R. Sabbatini, April 2004.
- 334 “Is inflation persistence intrinsic in industrial economies?” by A. T. Levin and J. M. Piger, April 2004.
- 335 “Has eura-area inflation persistence changed over time?” by G. O'Reilly and K. Whelan, April 2004.
- 336 “The great inflation of the 1970s” by F. Collard and H. Dellas, April 2004.
- 337 “The decline of activist stabilization policy: Natural rate misperceptions, learning and expectations” by
A. Orphanides and J. C. Williams, April 2004.

- 338 “The optimal degree of discretion in monetary policy” by S. Athey, A. Atkeson and P. J. Kehoe, April 2004.
- 339 “Understanding the effects of government spending on consumption” by J. Galí, J. D. López-Salido and J. Vallés, April 2004.
- 340 “Indeterminacy with inflation-forecast-based rules in a two-bloc model” by N. Batini, P. Levine and J. Pearlman, April 2004.
- 341 “Benefits and spillovers of greater competition in Europe: A macroeconomic assessment” by T. Bayoumi, D. Laxton and P. Pesenti, April 2004.
- 342 “Equal size, equal role? Interest rate interdependence between the euro area and the United States” by M. Ehrmann and M. Fratzscher, April 2004.
- 343 “Monetary discretion, pricing complementarity and dynamic multiple equilibria” by R. G. King and A. L. Wolman, April 2004.
- 344 “Ramsey monetary policy and international relative prices” by E. Faia and T. Monacelli, April 2004.
- 345 “Optimal monetary and fiscal policy: A linear-quadratic approach” by P. Benigno and M. Woodford, April 2004.
- 346 “Perpetual youth and endogenous labour supply: a problem and a possible solution” by G. Ascari and N. Rankin, April 2004.
- 347 “Firms’ investment decisions in response to demand and price uncertainty” by C. Fuss and P. Vermeulen, April 2004.
- 348 “Financial openness and growth: Short-run gain, long-run pain?” by M. Fratzscher and M. Bussiere, April 2004.
- 349 “Estimating the rank of the spectral density matrix” by G. Camba-Mendez and G. Kapetanios, April 2004.
- 350 “Exchange-rate policy and the zero bound on nominal interest rates” by G. Camba-Mendez and G. Kapetanios, April 2004.
- 351 “Interest rate determination in the interbank market” by V. Gaspar, G. P. Quirós and H. R. Mendizábal, April 2004.
- 352 “Forecasting inflation with thick models and neural networks” by P. McNelis and P. McAdam, April 2004.
- 353 “Towards the estimation of equilibrium exchange rates for CEE acceding countries: methodological issues and a panel cointegration perspective” by F. Maeso-Fernandez, C. Osbat and B. Schnatz, April 2004.
- 354 “Taking stock: monetary policy transmission to equity markets” by M. Ehrmann and M. Fratzscher, May 2004.
- 355 “Production interdependence and welfare” by K. X. D. Huang and Z. Liu, May 2004.

- 356 “Developing a euro area accounting matrix: issues and applications” by T. Jellema, S. Keuning, P. McAdam and R. Mink, May 2004.
- 357 “Seasonal adjustment and the detection of business cycle phases” by A. M. Mir, and D. R. Osborn, May 2004.

