



EUROPEAN CENTRAL BANK
EUROSYSTEM

Working Paper Series

Fédéric Holm-Hadulla, Kirstin Hubrich

Macroeconomic implications of
oil price fluctuations:
a regime-switching framework for the
euro area

No 2119 / December 2017

Abstract

We investigate whether the response of the macro-economy to oil price shocks undergoes episodic changes. Employing a regime-switching vector autoregressive model we identify two regimes that are characterized by qualitatively different patterns in economic activity and inflation following oil price shocks in the euro area. In the *normal regime*, oil price shocks trigger only limited and short-lived adjustments in these variables. In the *adverse regime*, by contrast, oil price shocks are followed by sizeable and sustained macroeconomic fluctuations, with inflation and economic activity moving in the same direction as the oil price. The responses of inflation expectations and wage growth point to second-round effects as a potential driver of the dynamics characterising the adverse regime. The systematic response of monetary policy works against such second-round effects in the adverse regime but is insufficient to fully offset them. The model also delivers (conditional) probabilities for being (staying) in either regime, which may help interpret oil price fluctuations – and inform deliberations on the adequate policy response – in real-time.

JEL Classification: E31, E52, C32

Keywords: regime switching models, time-varying transition probabilities, oil prices, inflation expectations, inflation

Non-technical summary

How has the sharp oil price decline since mid-2014 affected macroeconomic prospects in the euro area? Like in previous episodes of major oil price fluctuations, this question has generated substantial debate – and strongly divergent assessments – in the economics profession. Adopting a benign view, several observers have argued that the lower oil price will support the economic recovery by raising real disposable incomes and profits of euro area households and firms. Others, instead, have cautioned that the oil price slump may become entrenched in inflation expectations, thus leading to second-round effects that reinforce the prevailing disinflationary pressures and potentially dampen the recovery.

From a monetary policy perspective, it is crucial to establish the relative merit of these different assessments. Judging by public statements by policy-makers in various jurisdictions, central banks would typically consider changing the monetary policy stance only in the latter scenario, in which oil price fluctuations feed through to inflation expectations, hence risking to exert durable effects on actual inflation dynamics. By contrast, absent such second-round effects, central banks would tend to preserve the prevailing monetary policy stance, thus ‘looking through’ the oil price fluctuations.

To operationalise this nuanced reaction function, central banks in turn have to take a stand on which scenario they consider more likely to prevail; and, since the occurrence of second-round effects may be an episodic phenomenon, they may have to update this assessment on a regular basis.

The aim of the current paper is to examine whether the notion of episodic changes in the macroeconomic implications of oil price fluctuations finds empirical support in the euro area context. To this end, we use a regime-switching vector autoregression (RS-VAR) model, estimated with Bayesian methods, that allows for time-variation in model coefficients, thus helping to detect regime-dependent differences in how economic activity and inflation evolve in the aftermath of oil price shocks (henceforth referred to as *coefficient switching*). The model also allows for time-variation in shock variances (*variance switching*) since regime-dependent differences in macroeconomic dynamics might as well result from sequences of unusual shocks in certain episodes.

In terms of methodology, we employ the regime-switching VAR model developed in Hubrich, Waggoner and Zha (2016). In contrast to earlier regime-switching VAR models, as proposed in Sims et al. (2008), this method includes time-varying transition probabilities

that depend on the state of the economy. Moreover, it is, to our knowledge, the first to incorporate time-varying transition probabilities in a VAR framework (rather than in single equation regression models).

The analysis provides evidence of two regimes that are characterized by qualitatively different patterns in economic activity and inflation following oil price shocks. In the *normal regime*, downward oil price shocks, for instance, are followed by only transitory and moderate declines in inflation and a small increase in economic activity, which may attenuate the initial disinflationary effect of the oil-price declines. In the *adverse regime*, by contrast, the responses of economic activity and inflation are much more sizeable and sustained than in the normal regime and both variables move in the same direction as the oil price shock. Market-based measures of inflation expectations, which may act as a symptom for the presence of second-round effects, also show striking differences across regimes. While, in both regimes, inflation expectations initially adjust in the same direction as the oil price shock, this impact is again moderate and transitory in the normal regime whereas it is pronounced and persistent in the adverse regime. Overall, the dynamics characterising the adverse regime are consistent with the patterns to be expected in the presence of second-round effects that may amplify and prolong the impact of an oil price shock on inflation. The systematic response of monetary policy generally works against such second-round effects in the adverse regime but is insufficient to fully offset them.

Model extensions point to wage-price spirals as an important channel by which second-round effects may operate. By contrast, we do not find support for an interest-rate channel, which would operate if falling inflation expectations put upward pressure on long-term real rates and hence dampened real activity in the adverse regime.

The analysis also delivers time-varying probabilities that allow us to form a view on whether the euro area economy is more likely to be in one regime or the other at a given point in time, as well as the conditional probability of staying in that regime. Since this information is available at regular (monthly) frequency, the model may provide information supporting the deliberations on the adequate response to observed oil price fluctuations in real-time. These features also help us comment on past episodes, including the oil-price declines observed since mid-2014 until early-2015. Our model indicates that the euro area economy was in an adverse regime at that time and hence favours a pessimistic assessment of their implications for the strength of the euro area economy.

Cheaper oil is a rare piece of good news for (...) the euro currency area, since [it] should boost the spending power of Europe's consumers (...) amid the eurozone's long slump.

(The Wall Street Journal, November 14, 2014)

[A] danger [of the oil-price slump] is that an even deeper dip in inflation (...) may have an unwelcome second-round effect by dragging down inflation expectations.

(The Economist, December 4, 2014)

1 Introduction

How has the sharp oil price decline since mid-2014 affected macroeconomic prospects in the euro area? Like in previous episodes of major oil price fluctuations, this question has generated substantial debate – and strongly divergent assessments – in the economics profession. Adopting a benign view, several observers have argued that the lower oil price will support the economic recovery by raising real disposable incomes and profits of euro area households and firms. Others, instead, have cautioned that the oil price slump may become entrenched in inflation expectations, thus leading to second-round effects that reinforce the prevailing disinflationary pressures and potentially dampen the recovery.¹

From a monetary policy perspective, it is crucial to establish the relative merit of these different assessments. Judging by public statements by policy-makers in various jurisdictions, central banks would typically consider changing the monetary policy stance only in the latter scenario, in which oil price fluctuations feed through to inflation expectations, hence risking to exert durable effects on actual inflation dynamics. By contrast, absent such second-round

¹See, for instance, Mohaddes and Pesaran (2016) for a benign interpretation and Obstfeld et al. (2016) for an adverse interpretation of this recent collapse in oil prices. An analogous debate took place in the context of the steep upward trend in oil prices starting in end-2003 and accelerating from early-2007 to mid-2008, when some observers expressed concerns that second-round effects may lead to sustained inflationary pressures, whereas others contested this claim. For a summary of the debate at that time, see Hannon (2008).

effects, central banks would tend to preserve the prevailing monetary policy stance, thus ‘looking through’ the oil price fluctuations.²

To operationalise this nuanced reaction function, central banks in turn have to take a stand on which scenario they consider more likely to prevail; and, since the occurrence of second-round effects may be an episodic phenomenon, they may have to update this assessment on a regular basis. At the same time, a large body of literature has shown that policy mistakes in either direction (*i.e.* overly activist monetary policy responses to oil price fluctuations that prove transient, as well as overly inertial monetary policy that allows inflation expectations to become unanchored) may severely hamper macroeconomic performance.³

The aim of the current paper is to examine whether the notion of episodic changes in the macroeconomic implications of oil price fluctuations finds empirical support in the euro area context. To this end, we use a regime-switching vector autoregression (RS-VAR) model, estimated with Bayesian methods, that allows for time-variation in model coefficients, thus helping to detect regime-dependent differences in how economic activity and inflation evolve in the aftermath of oil price shocks (henceforth referred to as *coefficient switching*). The model also allows for time-variation in shock variances (*variance switching*) since regime-dependent differences in macroeconomic dynamics could also result from sequences of unusual shocks in certain episodes. Accounting for variance switching helps avoid that these differences are wrongly attributed to coefficient switching, which is the key hypothesis in the debate on second-round effects.⁴

In terms of methodology, we employ the regime-switching VAR model developed in Hubrich, Waggoner and Zha (2016). In contrast to earlier regime-switching VAR models,

²For statements that could be interpreted in this direction, see for instance ECB President Mario Draghi before the European Parliament on June 23, 2011: *In principle, if commodity price changes are of a temporary nature, one can look through the volatility in inflation triggered by their first-round effects. However, the risk of second round effects must be contrasted (...) to prevent that they have a lasting impact on medium-term inflation expectations. (...) In such cases, an adjustment of the monetary policy stance would be required to preserve price stability and keep inflation expectations well-anchored.* A similar distinction emerges from Chair Yellen’s assessment in the December 16, 2015 press conference: *For a number of years between 2004 and 2008, we had a series of increases in oil prices that (...) raised inflation (...) and we judged those increases to be transitory as well and looked through them. We do monitor inflation expectations very carefully. If we saw in a meaningful way that inflation expectations were either moving up in a way that made them seem unanchored or down, that would be of concern.*

³See, for instance, Bernanke et al. (1997); Barsky and Kilian (2001); Bodenstein et al. (2008); Nakov and Pescatori (2010); Kormilitsina (2011); Bodenstein et al. (2012); and Natal (2012).

⁴The importance of time-variation in the variance of key macroeconomic aggregates has most clearly emerged from the analysis of competing explanations for the steep decline in inflation and output volatility since the mid-1980s (often referred to as the Great Moderation); see, for instance, Primiceri (2005); Sims and Zha (2006).

as proposed in Sims et al. (2008) and employed in Hubrich and Tetlow (2015), this method includes time-varying transition probabilities that depend on the state of the economy. Moreover, it is, to our knowledge, the first to incorporate time-varying transition probabilities in a VAR framework (rather than in single equation regression models; see section 2 for further discussion).

The analysis provides evidence of two regimes that are characterized by qualitatively different patterns in economic activity and inflation following oil price shocks. In the *normal regime*, downward oil price shocks, for instance, are followed by only transitory and moderate declines in inflation and a small increase in economic activity, which may attenuate the initial disinflationary effect of the oil-price declines. In the *adverse regime*, by contrast, the responses of economic activity and inflation are much more sizeable and sustained than in the normal regime and both variables move in the same direction as the oil price shock. Market-based measures of inflation expectations, which may act as a symptom for the presence of second-round effects, also show striking differences across regimes. While, in both regimes, inflation expectations initially adjust in the same direction as the oil price shock, this impact is again moderate and transitory in the normal regime whereas it is pronounced and persistent in the adverse regime. Overall, the dynamics characterising the adverse regime are consistent with the patterns to be expected in the presence of second-round effects that may amplify and prolong the impact of an oil price shock on inflation. The systematic response of monetary policy generally works against such second-round effects in the adverse regime but is insufficient to fully offset them.

Model extensions point to wage-price spirals, of the type described in Hofmann et al. (2012), as an important channel by which second-round effects may operate. By contrast, we do not find support for a interest-rate channel (proposed, e.g., by Obstfeld et al. (2016)), which would operate if falling inflation expectations put upward pressure on long-term real rates and hence dampened real activity in the adverse regime (with analogous dynamics in response to rising inflation expectations; see also Bodenstein et al. (2013)).

The analysis also delivers time-varying probabilities that allow us to form a view on whether the euro area economy is more likely to be in one regime or the other at a given point in time, as well as the conditional probability of staying in that regime, provided it has previ-

ously prevailed. Since this information is available at regular (monthly) frequency, the model may provide information supporting the deliberations on the adequate response to observed oil price fluctuations in real-time. These features also help us comment on past episodes and, in particular, the recent debate on the macroeconomic implications of the oil-price declines observed since mid-2014 until early-2015. Our model indicates that the euro area economy was in an adverse regime at that time and hence favours a pessimistic assessment of their implications for the strength of the euro area economy. Indeed, counterfactual experiments show that economic activity, inflation and inflation expectations over the second half of 2014 were much more sluggish than we would have expected, had the euro area not been mired in the adverse regime. These findings are robust to several modifications in the model set-up and sample periods considered in the analysis.

Relation to the literature. Our paper contributes to an active literature aiming to shed light on how the implications of oil price shocks may differ depending on macroeconomic circumstances. One strand of this literature focuses on the *sources* of the underlying oil price shocks as a determinant of its macroeconomic implications. To this end, several papers have used structural VARs to model the global crude oil market and found that the macroeconomic consequences of oil price fluctuations differ depending on whether they originate from an oil supply shock, an increase in aggregate demand, or an increase in the precautionary demand for oil (see, for instance, Kilian (2008, 2009), Peersmann and Van Robays (2009), Lippi and Nobili (2012), Baumeister and Peersman (2013), Kilian and Murphy (2014), Baumeister and Hamilton (2015), and Caldara et al. (2016)). A second strand of this literature focuses on differences in the *transmission* of oil price shocks, either by comparing pre-defined historical episodes (as in Blanchard and Galí (2007) and Nakov and Pescatori (2010)); or by developing models that explicitly allow for non-linearities (as in Hamilton (2003)) and time-variation in the impact of oil shocks (as in Hahn and Mestre (2011), Van Robays (2012), Baumeister and Peersman (2013), and Leduc et al. (2016)). A third strand relevant to our analysis evaluates how central banks respond to oil price shocks, *inter alia* to explore whether changes in the systematic monetary policy response may explain the decline in the reduced-form relationship between oil prices and other macroeconomic aggregates during the Great Moderation.⁵

⁵In addition to some of the papers quoted in footnote 3 see, for instance, Bernanke et al. (2004), Hamilton and Herrera (2004) Herrera and Pesavento (2009), and Bjornland et al. (2016).

Our paper is closest to the second strand of the literature in that it also explicitly models time-variation in the effect of oil price shocks. But, to our knowledge, it is the first paper to: *(i)* account for such time variation in the euro area context; *(ii)* grant an explicit role to inflation expectations in the propagation of oil price shocks; and *(iii)* use a regime-switching model with endogenous switching to analyse oil price shocks.

Moreover, our model provides a new perspective on how oil price fluctuations enter monetary policy reaction functions – an issue that has been subject to a long-standing debate in the related literature. Initial contributions to this literature, focusing on the US context, point to a pronounced systematic monetary policy response to oil price shocks that counteracts their impact on inflation (see Bernanke et al. (1997)). Subsequent analysis instead has favoured a more nuanced reaction function that is conditional on the source of the shock; in particular, both VAR- and DSGE-based analyses, indicate that US monetary policy typically counteracts oil price shocks if they originate from global demand pressures, while adopting the opposite response to oil price shocks driven by supply disruptions (see, e.g., Kilian and Lewis (2011) and Bodenstein et al. (2012)). Our paper provides further evidence of such nuanced response, in that euro area monetary policy has tended to counteract oil price shocks only when they risk inducing second-round effects.

The remainder of the paper is organized as follows. Section 2 presents the methodology underlying the regime-switching model. Section 3 shows how the estimated economic adjustments in response to oil price shocks differ depending on the prevailing regime; it also reports how the estimated (conditional) probability of being (staying) in a certain regime has evolved since the start of the sample period; and it zooms in on the episode of collapsing oil prices since mid-2014, using counterfactual experiments to trace out how the euro area economy would have been expected to behave, if had it not switched to the adverse regime. Section 4 modifies and extends the baseline specification to explore different channels giving rise to differences across regimes and to expand the sample back to the 1970s for some of the estimations. Section 5 concludes.

2 Methodology

We employ the Regime-Switching Vectorautoregressive (RS-VAR) model with time-varying transition probabilities developed in Hubrich, Waggoner and Zha (2016).

Most of the methodological literature focusses on models with constant probabilities of Markov switching, including the seminal paper by Hamilton (1989), as well as the subsequent contributions by Chauvet (1998), Kim and Nelson (1999), Frühwirth-Schnatter (2004), Sims and Zha (2006), Sims et al. (2008) and Hubrich and Tetlow (2015).

In this paper, instead, we use a model that makes the transition probability dependent on the state of the economy, and therefore time-varying. This provides insights about how one or more variables in the system may affect the probability of the economy staying in one particular regime.

A parallel strand of the literature allows for time-variation in the probability of regime-switching but resorts to univariate or multivariate regression set-ups (see Filardo (1994), Diebold et al. (1994), Kim (2004), Kim et al. (2008), Amisano and Fagan (2013), Bazzi et al. (2014) as well as Chang et al. (2017)). In these papers, the probability of regime switching depends on certain variables of interest, but the regression set-ups do not permit feedback effects among the endogenous variables.

Building on and extending the framework presented in Sims et al. (2008) (SWZ08), the RS-VAR model by Hubrich et al. (2016) used in this paper combines these different strands of the literature. It allows for time-varying transition probabilities which depend on the state of the economy while modelling the interdependencies of the endogenous variables and imposing structural identifying assumptions in a regime switching structural vector autoregressive model with time-varying transition probabilities.

In this paper we employ this methodology to investigate the transmission of oil price shocks. Given our focus, we choose oil price changes as the driver of switching between different regimes. This facilitates relating the interpretation of the resulting regimes to oil price changes and their transmission through the macroeconomy.

2.1 Model

For $1 \leq t \leq T$, let y_t be an n -dimensional vector of endogenous variables, let z_t be an m -dimensional vector of exogenous variables, and let s_t be a discrete latent variable taking h distinct values. Let Θ be a vector of parameters controlling the distribution of y_t and let q be a vector of parameters controlling the process s_t . We will denote $\{y_1, \dots, y_t\}$ by Y_t , $\{z_1, \dots, z_t\}$ by Z_t , and $\{s_1, \dots, s_t\}$ by S_t . We assume the distribution of the exogenous variable z_t satisfies

$$p(z_t | s_t, Y_{t-1}, Z_{t-1}, \Theta, q) = p(z_t | Z_{t-1}). \quad (1)$$

In the ensuing analysis, we consider the time varying structural vector autoregression (RS-SVAR) defined by

$$A_0(s_t^c)y_t = A_+(s_t^c)x_t + \Xi^{-1}(s_t^v)\varepsilon_t,$$

where

$$x_t' = [y_{t-1}', \dots, y_{t-p}', 1],$$

$\Xi^{-1}(s_t^v)$ is a diagonal matrix with the standard errors of the shocks on the diagonal, ε_t is a standard normal vector of shocks, and s_t^c and s_t^v take on values in $\{1, \dots, h^c\}$ and $\{1, \dots, h^v\}$, respectively. In the notation of equation 1, $s_t = (s_t^c; s_t^v)$; the number of regimes for each process is $h = h^c, h^v$ and the regime changes in the coefficients are driven by the variable(s) influencing the probability of staying in a certain regime, while for the variance switching we assume a Markov process. The only exogenous variable is a constant, $z_t = 1$. In this paper, we allow for two coefficient regimes and for two shock-variance regimes. The vector Θ consists of the elements of $A_0(s_t^c)$, $A_+(s_t^c)$, and $\Xi(s_t^v)$.

2.2 The transition matrix

We denote the time-varying probability for switching from regime j to i , $p(s_{t+1} = i | s_t = j, Y_t, Z_t, \Theta, q)$, by $p_{i,j,t}$ and assume that the diagonal elements of the transition matrix, which give the time-varying persistence of the j^{th} regime, are of the form:

$$p_{j,j,t} = \frac{1}{1 + e^{-u_{j,t}}}$$

where

$$u_{j,t} = c_j + \gamma_j y_{t-k+1}.$$

and $y'_{t-k+1} = [y'_t, \dots, y'_{t-k+1}]$. With two regimes the off-diagonal elements will be of the form $(1 - p_{j,j,t})$. For more than two regimes we assume that the off-diagonal elements are Dirichlet-scaled. To achieve parsimony in the context of this complex specification, in our model we choose $k = 1$ and most of the elements of γ_j will be restricted to zero. The off-diagonal will be of the form:

$$p_{i,j,t} = (1 - p_{j,j,t}) \hat{p}_{i,j}$$

where $\sum_{i \neq j} \hat{p}_{i,j} = 1$. The prior on the parameters c_j and γ_j will be normal. In the general notation presented here, the probability of staying in a particular coefficient regime can depend on one or more (and possibly all) of the endogenous variables. In our model, we are interested in identifying different regimes for the transmission of oil price fluctuations. Therefore, we make the probability of staying in a particular coefficient regime dependent on oil price changes.

2.3 Estimation procedure, data and identification

We present our empirical results in terms of probabilities and impulse responses as estimates at the posterior mode. The posterior mode is computed via a blockwise BFGS optimization algorithm following SWZ08. It is well known that the quality of the estimation of the posterior mode depends on choosing good initial values for the algorithm. We used the Dynamic Striated Metropolis Hastings Sampler (Waggoner, Wong and Zha, 2015) to simulate the posterior distribution. We used a number of random draws as initial values for the optimization routine and chose the posterior model estimate that corresponded to the highest peak of all resulting local peaks. A by-product of this sampler is the computation of the marginal data density (MDD) that we use to compare the statistical fit of different alternative specifications.

The RS-VAR, in the baseline specification, includes: euro area industrial production as a measure of economic activity (which, in contrast to data on real GDP, is available at monthly frequency); the euro area Harmonized Index of Consumer Prices (HICP) as a measure of inflation; the Brent crude oil price (in current US-Dollars); the bilateral US-Dollar/Euro exchange

rate;⁶ 5-year/5-year break-even inflation rates as a market-based measure of long-term inflation expectations; and the 3-month EURIBOR as a measure for short-term interest rates.⁷

For identification, we apply a Cholesky decomposition. In our baseline specification, the variables are ordered as in the previous paragraph, which implies that we impose a zero restriction on the contemporaneous effect of oil price shocks on economic activity and inflation, while allowing the oil price to react to these variables immediately. This assumption is consistent with the approach adopted, for instance, in Bernanke et al. (1997) and Christiano et al. (1996) and may be motivated by the role of oil as a globally-traded financial asset whose price could, in principle, reflect any relevant information at high frequency, whereas adjustments in industrial production and overall price dynamics in the economy are likely to proceed at a more sluggish pace.⁸ At the same time, substantial parts of the related literature have adopted an alternative approach, treating oil prices as predetermined with respect to innovations in economic activity and inflation – an assumption that received empirical support from Kilian and Vega (2011).⁹

Against this background, we conducted two exercises to inform our modeling approach. First, in line with the related literature, we assessed the robustness of our model to an alternative ordering (with oil prices ranking first) and found our key findings to remain broadly unchanged.¹⁰ Second, we replicated the analysis in Kilian and Vega (2011) for the euro area sample considered in the current paper (see appendix A.1). Based on this analysis, we cannot rule out a contemporaneous effect of euro area macroeconomic news on oil prices over the sample period, thus arguing for the ordering adopted in the baseline specification as the more prudent approach in this context.

All models with regime switching either in the variances or the coefficients, or both, outperform the constant parameter model in terms of their marginal data densities; accordingly, we statistically reject the hypothesis that there are no non-linearities. The model with two

⁶One of the modifications of our baseline specification directly include oil prices in Euro and therefore drop the exchange rate from the system of equations.

⁷Industrial production, HICP and oil prices are included as year-on-year log-changes and the remaining variables are defined in percent.

⁸See Fratzscher et al. (2014).

⁹See, e.g., Davis and Haltiwanger (2001), Kiseok and Ni (2002), Leduc and Sill (2004), Blanchard and Galí (2007), Kilian and Lewis (2011) and Stock and Watson (2016).

¹⁰For example, Blanchard and Galí (2007) explore different orderings for the price of oil whereas Rotemberg and Woodford (1996) approach similar endogeneity concerns through sample splits.

variance and two coefficient regimes and constant transition probabilities is not significantly different from that same model with a time-varying transition matrix. This, in turn, implies that estimating additional parameters to allow for a time-varying probability of staying in a regime conditional on the state of the economy does not deteriorate the statistical fit, while providing further insights on the role of oil prices in determining economic dynamics.

The basic sample includes euro area data at monthly frequency over the period from February 2004 to January 2015 and, for some of the model extensions, the period from February 2004 to December 2015. The starting point of the sample is dictated by the availability of data on break-even inflation rates, which were not recorded in a consistent manner prior to February 2004 (nor were any other market-based measures of long-term inflation expectations and survey-based measures of inflation expectations are generally not available at monthly frequency).

The use of different sample end-dates is motivated by the evolution of the macroeconomic environment over recent years: to keep the baseline model tractable, it includes one short-term interest rate variable to capture the prevailing monetary policy stance; however, since mid-2014, and even more forcefully since end-January 2015, the ECB has started relying on asset purchase programmes to inject additional monetary policy accommodation, also reflecting an increasingly limited scope for further reductions in monetary policy interest rates. Since such central bank asset purchases directly intervene at the longer end of the yield curve, it appears more suitable to move to a broader specification in this environment, including also a longer-term interest rate to provide a more comprehensive characterization of the monetary policy stance.

For this reason, we first estimate the baseline model over a shorter sample period until end January 2015, before the ECB engaged in its expanded asset purchase programme (see section 3); second, we extend the model to also include data for the entire year 2015 and add a measure for the long-term interest rate as a further variable to the model (see section 4).

3 Regime-dependent effects of oil price shocks – baseline results

3.1 Differences in the growth- and inflation-response across regimes

The results from our baseline specification, shown in Figure 1, provide evidence of two regimes that are characterized by very different patterns in economic activity and inflation in response to the oil price shock. In the *normal regime* (slashed blue line) oil price shocks are followed by only small macroeconomic fluctuations. Inflation briefly declines after the shock, but only by a few basis points and the decline is fully reversed over the 24-months horizon for which impulse responses are shown. Inflation expectations follow a similar path as actual inflation, declining slightly after the shock but then recovering after a few months. Economic activity, in turn, increases slightly after the shock (in line with the benign interpretation of the oil price shock as supporting domestic demand).

In the *adverse regime* (solid red line), the oil price shock is slightly larger on impact (in both regimes, the shock is calculated as one standard deviation of the observed oil price moves over the time periods when the respective regime prevailed; accordingly, the difference in the magnitudes of the respective shocks indicates that oil prices tended to be somewhat more volatile in the adverse regime).¹¹ But the oil price then recovers more quickly than in the normal regime and, overall, does not display strikingly different patterns across regimes.

The impact of the shock on the other variables in the system does, however, differ in relevant ways. Instead of rising, as in the normal regime, economic activity in the adverse regime declines, hitting a trough after two quarters that is almost a percentage point below steady state and, despite some recovery, remains below steady state levels until the end of the horizon. In contrast to the normal regime, where economic activity acts as a stabiliser, this downward effect of the oil price shock on economic activity thus tends to amplify the disinflationary forces. As a consequence, actual HICP inflation undergoes a pronounced and persistent decline, settling almost 0.2 percentage points below steady state by the end of the horizon. This decline is also reflected in inflation expectations, which show a similarly modest drop on impact as in the normal regime, but then continue to drift down over the entire horizon.¹²

¹¹Impulse response functions that normalize the oil price decline on impact across regimes yield patterns that are very similar to those shown in Figure 1.

¹²One could argue that, under certain circumstances, an adjustment in inflation expectations as observed in the adverse regime may support the central bank in counteracting deviations of inflation rates from its preferred

Overall, the dynamics observed in the adverse regime are consistent with second-round effects exerting a protracted adverse impact on the economy that translates into declines in actual as well as expected inflation (further intuition and an assessment of the channels that may give rise to second-round effects are provided in the below Discussion and section 4).¹³

Failing to account for these regime-dependent dynamics may lead observers to miss important characteristics of the inflation process and economic activity in the aftermath of an oil price shock. This becomes clear when comparing the impulse response functions from a constant parameter VAR (also plotted in Figure 1 via the dotted black line) with those from the regime-switching model.¹⁴ In particular, restricting coefficients and variances to be constant may lead observers to underestimate the pronounced and persistent effects that oil price shocks may exert on economic activity and inflation in the adverse regime, while providing the wrong sign for output and inflation response in the normal regime. Both types of misjudgement may, in turn, contribute to policy mistakes in the response to oil price shocks.

The regime-switching framework also provides for a richer characterization of the monetary policy response to oil price shocks. The constant parameter VAR would point to a rather mechanical response by which monetary policy would always tend to accommodate the oil price shocks (via a monetary tightening following upward shocks and a loosening following downward shocks to the price of oil). Consistent with recent literature and with central bank communication, however, the RS-VAR instead favours a more nuanced assessment. In fact, monetary policy accommodates the oil price shock only in the adverse regime, *e.g.* leading to declining short-term interest rates in response to a drop in oil prices (although the resultant monetary loosening is not sufficient to offset the disinflationary forces). By contrast, in the normal regime, the short-term interest rate instead would tend to rise following the downward oil price shock, thus counteracting the ensuing boost in economic activity. These patterns echo

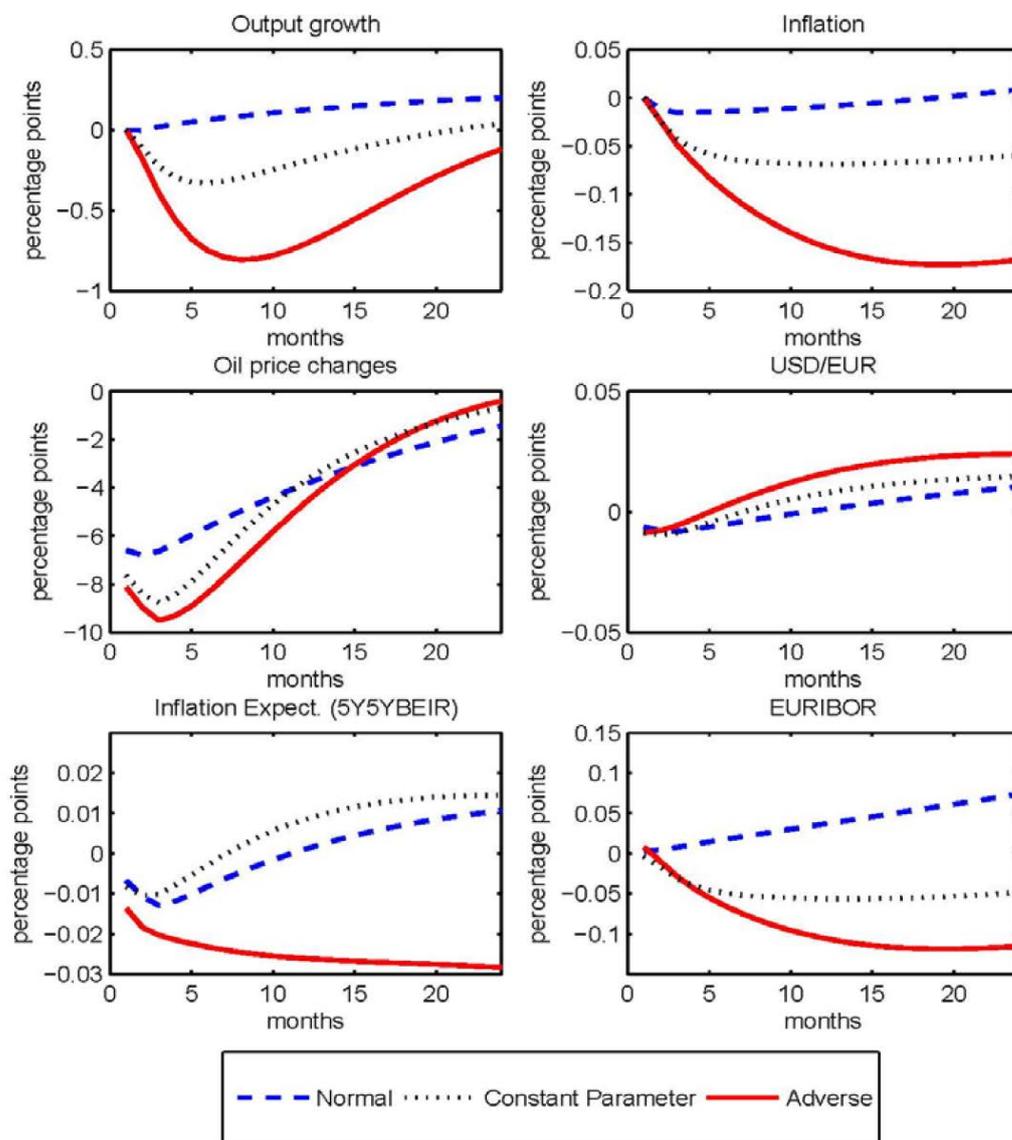
levels. For instance, if oil prices suddenly rise in conditions of below-target inflation rates, an upward adjustment in inflation expectations may accelerate the return to target. Still, we consistently refer to this regime as “adverse” based on the well-established notion that a de-anchoring of inflation expectations generally renders a central bank’s macroeconomic stabilization objectives harder to achieve; see Woodford (2007).

¹³Exchange rates do not play a major role in the adjustment dynamics, possibly reflecting the countervailing effects of the declining oil price, which would tend to induce an appreciation of the Euro versus the US-Dollar, and the weakening euro area economy, which would induce a depreciation.

¹⁴In terms of overall size, the impulse response of inflation in the constant parameter VAR is broadly comparable to the benchmark estimates presented in ECB (2006), as well as to the results from the estimated DSGE model by Forni et al. (2015). The former study points to a 0.1-0.2 percentage point inflation impact for a 10 percent change in oil prices and the latter finds a response that is at the lower end of this range.

the findings from recent literature aiming to distinguish monetary policy responses to different types of shocks (see section 3.3 for further discussion).

Figure 1: Impulse responses to oil price shock (one standard deviation shock), slashed blue: normal regime, solid red: adverse regime



3.2 Episodes of second-round effects

Figure 2 displays smoothed probabilities of being in a normal regime (with grey-shaded areas showing periods in which this probability exceeded 0.5) and the time-varying conditional probabilities of staying in a normal regime (depicted by the solid black line). We choose the threshold of 0.5 for presentational ease. At the same time, the estimated probabilities point

to a sharp distinction between regimes, with most of the regime-switches taking the form of jumps from probabilities close to zero to probabilities close to one (rather than of fluctuations around the 0.5 threshold).

The smoothed probabilities show that the euro area economy entered the adverse regime at various occasions since the start of the sample. Typically, regime-switches occurred after a sequence of pronounced, unidirectional oil price changes – for instance in the episodes of strong oil price declines and subsequent increases over the period 2008-2011 (for reference, see Figure 10 plotting the evolution of oil prices over the sample period). In August-2014, the euro area economy again switched to the adverse regime and, after a short period in the normal regime in the last quarter of 2014, returned to the adverse regime by the turn of the year.

Meanwhile, the conditional probability of staying in a normal regime (given the economy was in that regime), which is a function of the oil price evolution, declined steeply just before the economy transitioned to the adverse regime in late-2008. After a prolonged period with a high conditional probability of staying in the normal regime (which starts in end-2009 – before the smoothed probabilities actually indicate a return to the normal regime), this probability again fell steeply in the second half of 2014 and approached zero by January 2015.

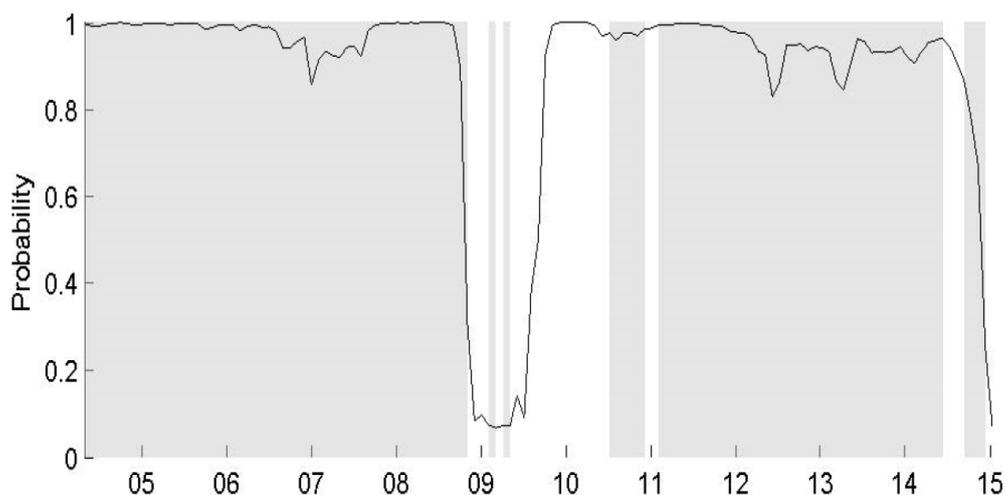
Accordingly, our model clearly assigns the period from August 2014 to January 2015, which had triggered intense debate among the policy observers (see section 1), to the adverse regime. In other words, the regime-switching framework indicates that the oil price declines observed over that period are likely to have reinforced the prevailing disinflationary pressures via second-round effects, thus favoring the more pessimistic assessment of that episode.

Since the model allows for coefficient- *and* variance-switching, it also yields insights as to whether the different types of regimes tend to coincide. To this end, Figure 3 presents the probability for the economy to be in a high-variance regime over the sample period.¹⁵ Overall, the comparison to Figure 2 indicates that variance-switching occurs more frequently than coefficient-switching and there is limited synchronicity between the different types of regimes. For instance, several switches to the high-variance regime arise in the pre-2008 period, when the normal coefficient regime seems to have prevailed; and during the adverse coefficient-regime periods, no clear pattern emerges with regard to variance-regimes. Accordingly, there

¹⁵Since we do not assume a time-varying transition matrix for the variance regimes (with the switching instead being driven by a Markov process), Figure 3 only contains one set of probabilities.

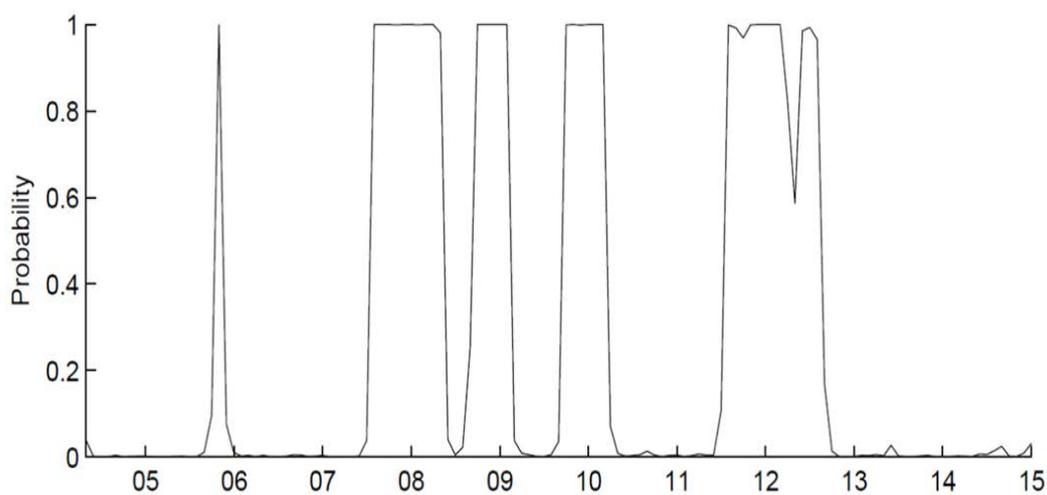
is little indication that adverse regimes tend to be particularly likely in times of high volatility, or vice versa.

Figure 2: Time-varying probability for being in a normal regime (grey-shaded area) and conditional probability of staying in that regime (black line)



Note: The ticks on the x-axis refer to the beginning of the respective year (with “05” indicating the year 2005 etc.)

Figure 3: Probability of high-variance regime



Note: The ticks on the x-axis refer to the beginning of the respective year (with “05” indicating the year 2005 etc.)

3.3 Regime-dependent transmission versus different sources oil price shocks

It is interesting to relate these findings to the strand of the literature that emphasises different sources of oil price shocks as a key determinant of their macroeconomic effects. As we will discuss, when comparing the impulse response functions between the normal and the adverse regimes, important parallels emerge with that literature. At the same time, these parallels remain partial and the time periods assigned to the different regimes in our framework display only limited overlap with the demand- versus supply-shock dichotomy from the related literature.

Parallels. In terms of the impulse response functions, the macroeconomic adjustments to oil price fluctuations in the normal regime resemble the dynamics one may expect after a genuine supply-driven shock: in response to falling oil prices, for instance, euro area economic activity temporarily strengthens, consistent with increased disposable income for euro area households and reduced energy input costs for euro area producers; meanwhile, the oil price decline initially exerts some moderate direct downward pressure on inflation, probably reflecting the role of energy in the consumption basket and some pass-through of lower input costs to consumer prices, but the decline is quickly offset by the uptick in activity. Consistent with the short-lived inflation response, inflation expectations quickly return to steady state.

In the adverse regime, by contrast, the co-movement patterns are closer to what may be expected in response to a global demand-driven oil price shock. Again resorting to the example of falling oil prices, the concomitant protracted slump in economic activity, which materialises despite an initial beneficial effect on euro area disposable incomes and input costs, is consistent with a situation in which the oil price decline signals a downgrade in broader economic prospects, inducing firms and households to revise down their expectations for future inflation and, via this channel, providing an incentive to postpone nominal spending decisions. This inter-temporal substitution, in turn, further depresses domestic activity and reinforces the downward pressure on inflation, thus giving rise to second-round effects.

The different monetary policy responses across regimes further strengthen these parallels. For instance, using a structural VAR, Kilian and Lewis (2011) find that US monetary policy tended to counteract only demand-driven oil price shocks, which the authors rationalise with the intention to preempt wage-price spirals, whereas supply-induced oil price shocks trigger a

monetary policy response in the opposite direction. This pattern, which has also found support in DSGE-frameworks (see, e.g., Natal (2012) and Bodenstein et al. (2012)), again conforms with the short-term interest rate dynamics we detect in the different regimes, with the normal regime mapping into a supply-shock and the adverse regime into a demand-shock scenario.

Differences. While these parallels are intuitively plausible, the overlap between the different approaches remains partial. First, the regime-switching framework produces a sharper distinction in the estimated impulse response functions between the different regimes than what the related literature finds for different types of shocks in the euro area context. For instance, Peersmann and Van Robays (2009) find a quantitatively similar, positive, impact of oil price increases on euro area inflation, independently of whether it is driven by a supply or demand shock (in contrast to the qualitatively different patterns observed in Figure 1).

Second, and perhaps most importantly, the periods assigned to the adverse regime in our model were not characterised by a particular dominance of demand shocks, according to the studies making this distinction. For instance, the shock decompositions in Baumeister and Hamilton (2015) indeed attribute an important part of the oil price collapse at the onset of the ‘Great Recession’ in 2008 to global economic activity and oil demand shocks. While our model also assigns the euro area economy to the adverse regime in this episode, this regime persists well beyond the period over which Baumeister and Hamilton (2015) see global economic activity and demand shocks to dominate; For instance, the shock decompositions in Baumeister and Hamilton (2015) indicate that also non-demand related shocks made a sizeable contribution to the oil price collapse at the onset of the ‘Great Recession’ in 2008 – *i.e.* the first period assigned to the adverse regime in our model; moreover, according to our model, the adverse regime, with two brief interruptions, persisted until mid-2010 – a period over which global activity and demand shocks were fairly mute according to Baumeister and Hamilton (2015); our model also sees the euro area economy re-enter the adverse regime around end-2010 to early-2011, when Baumeister and Hamilton (2015) identify a strong supply-contribution to the oil price increase that materialised around this time; and a similar pattern emerges for the oil price drop in the second half of 2014, when our model points to a strong decline in the probability of staying in the normal regime (see also Baumeister and Kilian (2016) for an in-depth analysis of this period).

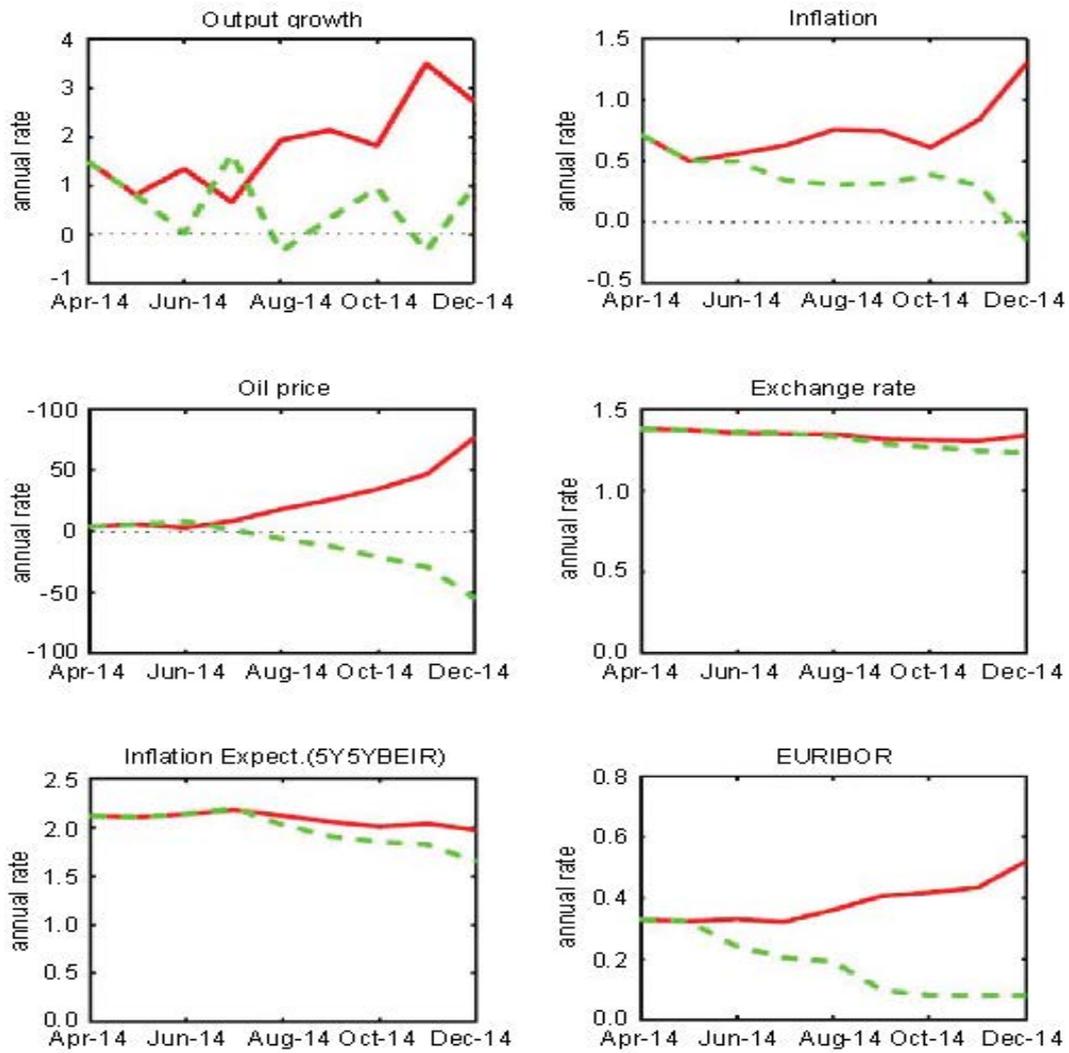
Accordingly, rather than just capturing the evolution of different types of shocks through a different lens, the regime-switching perspective constitutes an interesting complement to the existing literature that disentangles supply- and demand-factors causing oil price changes. In particular, by allowing for different propagation patterns of oil price shocks at different points in time and determining the probability of being in a specific regime, the model may add to policy-makers' information sets when deciding on the appropriate response to an observed oil price fluctuation.

3.4 Consequences of a regime-switch

As a final exercise with our baseline model, Figure 4 presents a counterfactual experiment that allows us to gain further insight into the contentious episode starting in the second half of 2014 and to illustrate how the dynamics of the system differ across regimes. The basic set-up of this counterfactual experiment is to trace out how the different variables would have evolved if the euro area economy had not switched to the adverse regime by August 2014.

The counterfactual points to striking differences in economic outcomes across regimes. Economic activity and inflation would have been more than 1.5 percentage point higher than observed by the end of 2014, had the normal rather than the adverse regime prevailed at that time. Moreover, inflation expectations would have continued to hover around 2%, rather than following a marked downward path to a new historical minimum of 1.6% by the end of 2014 – in contrast to the early years of the crisis, when inflation expectations faced upward pressure from rising oil prices thus contributing to the “missing disinflation” in 2009-2011 (see Coibion and Gorodnichenko (2015)). Interestingly, also the oil price would have staged a dynamic recovery instead of the actual steep decline. This indicates that the weakness in euro area economic activity and inflation from mid-2014 to early-2015 has been one of the drivers of the global oil price collapse over the second half of 2014.

Figure 4: Counterfactual experiment: normal instead of adverse regime (green dashed line: actual data, red solid line: counterfactual path)



4 Extensions

While the adverse-regime dynamics in the baseline specification are consistent with the presence of second-round effects, further analysis is necessary to identify potential channels by which these effects may operate. To explore this issue and to test the robustness of our key findings, we modify and extend our model in several directions. First, we estimate a more parsimonious specification that directly expresses oil prices in euro and omits the exchange rate from the model; the resultant reduction in the size of the VAR is convenient in view of the additional variables that are later added to the model to study different channels and, from an economic perspective, may be motivated by the fact that ultimately it is the oil price in euro that matters for euro area consumers and firms. Second, we augment the model with measures for euro area wage growth and the long-term real interest rate, so as to study two alternative channels that the related literature has proposed as giving rise to second-round effects. Third, we re-run the regime-switching model on a different euro area data set that extends significantly further back in time (ranging back to the 1970s), while requiring us to move to quarterly frequency and to exclude inflation expectations from the set of variables.

4.1 Oil price in euro – a more parsimonious specification

As apparent from Figures 5 and 6, the specification with oil prices in euro confirms our key results (see appendix A.2). In particular, the impulse response functions show nearly identical patterns for economic activity, inflation, and inflation expectations as in the baseline, and the differences across regimes are similarly pronounced. Also, the assignment of periods to the two regimes is broadly unchanged, albeit with somewhat more frequent switches to the adverse regime over the period 2010-2012 than in the baseline specification. The period around the turn of the year 2014 to 2015 is again assigned to the adverse regime and the drop in the conditional probability of staying in the normal regime in the second half of 2014 is also confirmed by this alternative specification. Accordingly, this more parsimonious model is a suitable substitute to the baseline when adding further variables, and in particular wages, to study the channels giving rise to second-round effects.

4.2 Wage-price spirals and the real interest rate channel

As explained by Hofmann et al. (2012), a wage-channel for second-round effects may operate if nominal wage growth rises (falls) in response to observed increases (declines) in the inflation rate, thereby amplifying the initial (dis-)inflationary push via a mutually reinforcing wage-price spiral (see p. 769). As a source for such wage-price spirals, the authors emphasise the existence of explicit or implicit wage indexation. Since wage indexation is determined by (typically slow-moving) changes in labour market institutions, it does not appear as a likely candidate for the relatively frequent regime-switches reported in Figures 2 and 6 of the current paper, which we instead rationalized with adjustments in inflation expectations. However, following a standard hybrid Phillips curve specification (see Galí and Gertler (1999) and Blanchard and Riggi (2013)), wage indexation might operate via analogous economic mechanisms as a de-anchoring of inflation expectations: like wage indexation, a de-anchoring of inflation expectations raises the weight on the backward-looking component of the hybrid Phillips curve, thus facilitating a similar type of wage-price spiral.¹⁶

The real interest rate channel, emphasised for instance by Obstfeld et al. (2016), would operate if the fall in inflation expectations triggered by the oil price decline were to induce an effective tightening in financial conditions which in turn dampens aggregate demand – a scenario that is particularly likely if monetary policy is constrained in its ability to inject additional accommodation to counteract the upward pressure on long-term real interest rates and arrest the decline in inflation expectations. As a consequence, even an initially benign oil-price decline may exert a negative effect on economic activity.

The evidence from our model confirms wages as a channel for second-round effects, whereas real interest rates do not emerge as a relevant factor. Figure 7 adds the annual growth rate of nominal negotiated wages in the euro area to the previous specification. Consistent with our earlier findings, economic activity, inflation and inflation expectations, in the adverse regime, follow oil prices on their way down. In the normal regime, by contrast, these variables show only a very small response and an opposite sign for economic activity (compared to earlier specifications there is a more dynamic recovery in the oil price, which also pulls up actual

¹⁶Formally, this becomes apparent from a comparison of the γ -parameters in equation 7 in Hofmann et al. (2012) and in equation 24 in Galí and Gertler (1999).

and expected inflation). As would be expected in the presence of wage-based second-round effects, wages – after an initial uptick – enter on a lasting downward path in the adverse regime whereas no such dynamics materialise in the normal regime.

Figure 8 extends the baseline model with a measure of the real long-term interest rate, defined as the 10-year sovereign bond yields of euro area countries, weighted by their respective shares in euro-area GDP and deflated with market based inflation expectations over the same maturity. It also extends the sample period until end-2015 since the inclusion of a long-term interest rate variable renders the model better suited to account for the effects of the ECB's expanded asset purchase programme adopted in January 2015 (see section 2.3).¹⁷ The impulse response functions again confirm the broad findings from our baseline specification. But, contrary to the hypothesis of a real-interest rate channel, which would require an increase in real rates in the adverse regime, the real rate falls. A possible interpretation is that monetary policy, through forward guidance and non-standard measures that directly intervene at the longer end of the yield curve, has been able to overcompensate the upward pressure on real rates triggered by the decline in inflation expectations, even as the scope for short-term interest rate cuts became limited.

4.3 Expanding the sample back to the 1970's

The use of inflation expectations, for which the data start only in 2004, has implied an exclusion of several interesting oil-price episodes from the scope of the analysis. As a final exercise, we address this limitation by re-running the RS-VAR on the Area-wide Model database, which provides quarterly data back to the 1970s (see Fagan et al. (2005)). In analogy to the baseline, the variables in the model include real GDP growth, HICP inflation, the oil price change in USD, the bilateral US-Dollar/Euro exchange rate, nominal per capita wage growth and a nominal short-term interest rate. The sample begins in 1974, when crude oil prices started being determined on an integrated global oil market (see Kilian (2014)).

Based on this longer sample, we again find opposing signs for the response of economic activity in the normal versus the adverse regime, with real GDP falling in response to a decline in the price of oil. At the same time, the real GDP response is weaker than that of industrial pro-

¹⁷For sake of comparability, we used the same sample also for the specification including wages shown in Figure 7.

duction in the previous specifications, consistent with the inclusion of several less-responsive expenditure components, such as government consumption, in the real GDP variable. Meanwhile, the response in inflation also differs across regimes, but in a more nuanced manner than over the shorter sample: the dynamics in the adverse regime are characterised by an adjustment in the same direction as the oil price that is similar to, albeit somewhat flatter than, in the baseline. Following our previous reasoning, the positive co-movement of economic activity and inflation with the oil price points to the presence of second-round effects – an interpretation that is also supported by the corresponding decline in wage growth. In the normal regime, by contrast, the inflation response is initially very subdued, thus echoing the cross-regime differences found for the shorter sample. However, inflation then continues to proceed on a secular downward path that even crosses that from the adverse-regime after around one and a half years. Accordingly, the regimes differ in the pace with which inflation adjusts to the shock, rather than in the direction of the response (in contrast to the baseline where the responses differ in sign). Given qualitatively similar patterns in wage growth, the results indicate that, for the longer sample, second-round effects constitute a more pervasive phenomenon than in the post-2003 period considered in our earlier specifications; the latter, in turn, are likely to be more representative for current conditions given the sample is confined to relatively recent observations.

As an aside, the comparison of the wage response with that in Figure 7 points to an interesting parallel with recent findings by Peneva and Rudd (2015), who document a marked decline in the pass-through of labor costs to consumer price inflation since the mid-1960's. While being mute on the direction of causality, a qualitatively similar pattern emerges from our estimates, with a given wage response coinciding with a much more subdued inflation response in the more recent sample (Figure 7) than over the longer sample period (Figure 9).¹⁸

It is noteworthy, that the assignment of different time periods to the two regimes displays substantial overlap with our previous specifications for the parts of the sample that coincide. Again, the post-Lehman crisis episode, as well as the period of collapsing oil prices starting in second half of 2014 are assigned to the adverse regime. The respective regime episodes are more persistent, however, in that some of the brief switches after 2009 back to the normal

¹⁸In fact, the nominal-wage and inflation responses are sufficiently different to imply real-wage responses of opposite sign in the two sample periods.

regime from the shorter sample are not picked up by this specification.

Over the earlier (non-overlapping) sample, three additional episodes are assigned to the adverse regime, which may explain some of the differences in the impulse responses compared to the baseline (even though the important broad features carry over): the first in the mid-1970s, the second from the mid-1980s to the early 1990s, and the third from the late 1990s to the mid-2000s. All three episodes were characterised by a pronounced co-movement in oil and broader consumer price inflation.

In particular, the latter two episodes correspond, respectively, to: *(i)* the period of declining and converging euro area inflation rates in the run-up to the Maastricht Treaty; and *(ii)* the period of rising inflation in the years right after the introduction of the euro. Accordingly, our model would indicate that the pre-euro area disinflation process may have been supported by downward second-round effects induced by the oil price declines over that period; and that the concerns about upward second-round effects that motivated a sequence of ECB rate hikes in the early years of the euro area may have been warranted at that time.¹⁹

The first adverse-regime episode, in turn, coincides with the oil price hikes around the 1973 Yom Kippur War and the ensuing Arab oil embargo in 1973-1974. By contrast, the late-1970s to early-1980s – another period characterised by pronounced oil price spikes – are not assigned to the adverse regime, consistent with the relatively short-lived increase in inflation around that time, which was swiftly reversed with the onset of a global dis-inflationary phase over the 1980s.

5 Conclusions

In this paper, we have analysed whether the dynamics of euro area economic activity and inflation following oil price shocks undergoes episodic changes. To this end, we used a regime-switching VAR with endogenous switching following a novel approach developed by Hubrich et al. (2016). We find that oil price fluctuations are typically followed by limited adjustments in inflation and economic activity, a situation we refer to as a normal regime. Occasionally, however, the economy enters into an adverse regime where oil price shocks herald sizeable

¹⁹Already in January 2000, the introductory statement to the ECB press conference referred to second-round effects as a risk to price stability.

and sustained macroeconomic fluctuations, with inflation (actual and expected), as well as economic activity, moving in the same direction as the oil price shock. Overall, the dynamics observed in the adverse regime are consistent with the presence of second-round effects of oil price shocks on growth and inflation, with wage-price spirals acting as a key channel for the positive co-movement between these variables.

Zooming in on the episode of collapsing oil prices from mid-2014 to early-2015 – which generated a lively debate among central bank observers on whether or not it warranted a monetary policy response – our model indicates that the adverse regime is likely to have prevailed at that time. Accordingly, concerns of negative second-effects reinforcing the then-prevailing disinflationary pressures appeared warranted. In fact, according to our baseline estimates, economic activity and inflation would have been more than 1.5 percentage point higher than observed, had the euro area economy not been mired in an adverse regime at that time. Moreover, inflation expectations would have continued to hover around 2%, rather than following a marked downward path to a new historical minimum of 1.6% by the end of 2014.

References

- Amisano, G. and Fagan, G. (2013). Money growth and inflation: A regime switching approach, *Journal of International Money and Finance* **33**: 118–145.
- Barsky, R. B. and Kilian, L. (2001). Do we really know that oil caused the Great Stagflation? A monetary alternative, *NBER Macroeconomics Annual 2001* **16**: 137–183.
- Baumeister, C. and Hamilton, J. D. (2015). Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks, *Manuscript*.
- Baumeister, C. and Kilian, L. (2016). Understanding the decline in the price of oil since June 2014, *Journal of the Association of Environmental and Resource Economists* **3**(1): 131–158.
- Baumeister, C. and Peersman, G. (2013). Time-varying effects of oil supply shocks on the US economy, *American Economic Journal: Macroeconomics* **5**(4): 1–28.
- Bazzi, M., Blasques, F., Koopman, S. J. and Lucas, A. (2014). Time varying transition probabilities for Markov regime switching models, *Tinbergen Institute Discussion Papers 14-072/III*, Tinbergen Institute.
- Bernanke, B. S., Gertler, M. and Watson, M. (1997). Systematic monetary policy and the effects of oil price shocks, *Brookings Papers on Economic Activity* **1**: 91–157.
- Bernanke, B. S., Gertler, M. and Watson, M. W. (2004). Reply: Oil shocks and aggregate macroeconomic behavior: The role of monetary policy, *Journal of Money, Credit and Banking* **36**(2): 287–291.
- Bjornland, H. C., Larsen, V. H. and Maih, J. (2016). Oil and macroeconomic (in)stability, *Working Paper 2016/12*, Norges Bank.
- Blanchard, O. J. and Galí, J. (2007). The macroeconomic effects of oil shocks. Why are the 2000s so different from the 1970s?, *NBER Working Paper* (13368).
- Blanchard, O. J. and Riggi, M. (2013). Why are the 2000s so different from the 1970s? A

- structural interpretation of changes in the macroeconomic effects of oil prices, *Journal of the European Economic Association* **11**(5): 1032–1052.
- Bodenstein, M., Erceg, C. J. and Guerrieri, L. (2008). Optimal monetary policy with distinct core and headline inflation rates, *Journal of Monetary Economics* **55**: S18–33.
- Bodenstein, M., Guerrieri, L. and Gust, C. J. (2013). Oil shocks and the zero bound on nominal interest rates, *Journal of International Money and Finance* **32**: 941–967.
- Bodenstein, M., Guerrieri, L. and Kilian, L. (2012). Monetary policy responses to oil price fluctuations, *IMF Economic Review* **60**(4): 470–504.
- Caldara, D., Cavallo, M. and Iacoviello, M. (2016). Oil price elasticities and oil price fluctuations, *International Finance Discussion Papers 1173*, Board of Governors of the Federal Reserve System (U.S.).
- Chang, Y., Choi, Y. and Park, J. Y. (2017). A new approach to model regime switching, *Journal of Econometrics* **196**(1): 127 – 143.
- Chauvet, M. (1998). An econometric characterization of business cycle dynamics with factor structure and regime switching, *International Economic Review* **39**(4): 969–996.
- Christiano, L. J., Eichenbaum, M. and Evans, C. (1996). The effects of monetary policy shocks: Evidence from the flow of funds, *The Review of Economics and Statistics* **78**(1): 16–34.
- Coibion, O. and Gorodnichenko, Y. (2015). Is the Phillips curve alive and well after all? Inflation expectations and the missing disinflation, *American Economic Journal: Macroeconomics* **7**(1): 197–232.
- Davis, S. J. and Haltiwanger, J. (2001). Sectoral job creation and destruction responses to oil price changes, *Journal of Monetary Economics* **48**: 465–512.
- Diebold, F. X., Lee, J.-H. and Weinbach, G. C. (1994). Regime switching with time-varying transition probabilities, in C. Hargreaves (ed.), *Nonstationary Time Series Analysis and Cointegration*, Oxford University Press, pp. 283–302.

- ECB (2006). The pass-through of global energy prices to the euro area HICP energy components, *Monthly Bulletin Box 3 February 2006*, European Central Bank.
- Fagan, G., Henry, J. and Mestre, R. (2005). An area-wide model for the euro area, *Economic Modelling* **22**(1): 39 – 59.
- Filardo, A. J. (1994). Business-cycle phases and their transitional dynamics, *Journal of Business & Economic Statistics* **12**(3): 299–308.
- Forni, L., Gerali, A., Notarpietro, A. and Pisani, M. (2015). Euro area, oil and global shocks: An empirical model-based analysis, *Journal of Macroeconomics* **46**: 295 – 314.
- Fratzscher, M., Schneider, D. and Van Robays, I. (2014). Oil prices, exchange rates and asset prices, *ECB Working Paper* (1689).
- Frühwirth-Schnatter, S. (2004). Estimating marginal likelihoods for mixture and Markov switching models using bridge sampling techniques, *The Econometrics Journal* **7**(1): 143–167.
- Galí, J. and Gertler, M. (1999). Inflation dynamics: A structural econometric analysis, *Journal of Monetary Economics* **44**(2): 195 – 222.
- Hahn, E. and Mestre, R. (2011). The role of oil prices in the euro area economy since the 1970s, *Working Paper Series 1356*, European Central Bank.
- Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle, *Econometrica* **57**(2): 357–384.
- Hamilton, J. D. (2003). What is an oil shock?, *Journal of Econometrics* **113**: 363–398.
- Hamilton, J. D. and Herrera, A. M. (2004). Oil shocks and aggregate macroeconomic behavior: The role of monetary policy: Comment, *Journal of Money, Credit and Banking* **36**(2): 265–286.
- Hannon, P. (2008). BIS: No evidence of second-round effects from high oil prices, *The Wall Street Journal*.

- Herrera, A. M. and Pesavento, E. (2009). Oil price shocks, systematic monetary policy, and the Great Moderation, *Macroeconomic Dynamics* **13**(1): 107–137.
- Hofmann, B., Peersman, G. and Straub, R. (2012). Time variation in U.S. wage dynamics, *Journal of Monetary Economics* **59**(8): 769 – 783.
- Hubrich, K. and Tetlow, R. T. (2015). Financial stress and economic dynamics: The transmission of crises, *Journal of Monetary Economics* **70**(1): 100–115.
- Hubrich, K., Waggoner, D. and Zha, T. (2016). Monetary policy and financial stability: Unconventional monetary policy, leverage and financial stress. Manuscript, Federal Reserve Board.
- Kilian, L. (2008). The economic effects of energy price shocks, *Journal of Economic Literature* **46**(4): 871–909.
- Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market, *American Economic Review* **99**(3): 1053–1069.
- Kilian, L. (2014). Oil price shocks: Causes and consequences, *Annual Review of Resource Economics* **6**(1): 133–154.
- Kilian, L. and Lewis, L. T. (2011). Does the Fed respond to oil price shocks, *The Economic Journal* **121**: 1047–1072.
- Kilian, L. and Murphy, D. P. (2014). The role of inventories and speculative trading in the global market for crude oil, *Journal of Applied Econometrics* **29**(3): 454–478.
- Kilian, L. and Vega, C. (2011). Do energy prices respond to U.S. macroeconomic news? A test of the hypothesis of predetermined energy prices, *The Review of Economic and Statistics* **93**(2): 660–671.
- Kim, C.-J. (2004). Markov-switching models with endogenous explanatory variables, *Journal of Econometrics* **122**(1): 127–136.
- Kim, C.-J. and Nelson, C. R. (1999). Has the U.S. economy become more stable? A Bayesian

- approach based on a Markov-switching model of the business cycle, *The Review of Economics and Statistics* **81**(4): 608–616.
- Kim, C.-J., Piger, J. and Startz, R. (2008). Estimation of Markov regime-switching regression models with endogenous switching, *Journal of Econometrics* **143**(2): 263–273.
- Kiseok, L. and Ni, S. (2002). On the dynamic effects of oil price shocks: A study using industry level data, *Journal of Monetary Economics* **49**: 823–852.
- Kormilitsina, A. (2011). Oil price shocks and the optimality of monetary policy, *Review of Economic Dynamics* **14**(1): 199–223.
- Leamer, E. E. (1978). *Specification searches: Ad hoc inference with non-experimental data*, New York: John Wiley and Sons, Inc.
- Leduc, S., Moran, K. and Vigfusson, R. J. (2016). Learning in the oil futures markets: Evidence and macroeconomic implications, *International Finance Discussion Papers 1179*, Board of Governors of the Federal Reserve System (U.S.).
- Leduc, S. and Sill, K. (2004). A quantitative analysis of oil-price shocks, systematic monetary policy, and economic downturns, *Journal of Monetary Economics* **51**: 781–808.
- Lippi, F. and Nobili, A. (2012). Oil and the macroeconomy: A quantitative structural analysis, *Journal of the European Economic Association* **10**(5): 1059–1083.
- Mohaddes, K. and Pesaran, M. H. (2016). Oil prices and the global economy: Is it different this time around?, *Cambridge Working Paper Economics* (1640).
- Nakov, A. and Pescatori, A. (2010). Oil and the Great Moderation, *The Economic Journal* **120**(543): 137–183.
- Natal, J.-M. (2012). Monetary policy response to oil price shocks, *Journal of Money, Credit, and Banking* **44**(1): 53–101.
- Obstfeld, M., Milesi-Ferretti, G. M. and Arezki, R. (2016). Oil prices and the global economy: It's complicated, *iMFdirect blog posted on March 24* .

- Peersmann, G. and Van Robays, I. (2009). Oil and the euro area economy, *Economic Policy* **24**(60): 603–640.
- Peneva, E. V. and Rudd, J. B. (2015). The passthrough of labor costs to price inflation, *Finance and Economics Discussion Series 2015-42*, Board of Governors of the Federal Reserve System (U.S.).
- Rotemberg, J. J. and Woodford, M. (1996). Imperfect competition and the effects of energy price increases on economic activity, *Journal of Money, Credit and Banking* **28**(4): 549–577.
- Sims, C. A., Waggoner, D. F. and Zha, T. (2008). Methods for inference in large multiple-equation Markov-switching models, *Journal of Econometrics* **146**(2): 255–274.
- Sims, C. A. and Zha, T. (2006). Were there regime switches in U.S. monetary policy?, *American Economic Review* **96**(1): 54–81.
- Stock, J. and Watson, M. (2016). Dynamic factor models, factor-augmented vector autoregressions, and structural vector autoregressions in macroeconomics, in J. B. Taylor and H. Uhlig (eds), *Handbook of Macroeconomics*, Vol. 2, Elsevier, pp. 415 – 525.
- Van Robays, I. (2012). Macroeconomic uncertainty and the impact of oil shocks, *ECB Working Paper* (1479).
- Woodford, M. (2007). The case for forecast targeting as a monetary policy strategy, *Journal of Economic Perspectives* **21**(4): 3–24.

A Appendices

A.1 Response of oil prices to euro area macroeconomic news

The related literature is split on whether oil prices may be treated as predetermined with respect to other key macroeconomic aggregates (see section 2.3). This question is central to the appropriate ordering of variables in recursive identification schemes of the type considered in the current paper. Kilian and Vega (2011) propose a test of the hypothesis of predetermined energy prices which relies on regressions of WTI crude oil and US gasoline prices on the surprise component of a wide range of US macroeconomic indicators (defined as the difference between data releases and ex ante survey expectations for the respective variable). When accounting for potential distortions in test statistics due to “data mining”,²⁰ the estimates in Kilian and Vega (2011) fail to reject the null hypothesis that energy prices are contemporaneously unaffected by macroeconomic news, regardless of which indicator is included as explanatory variable; whether the regressions are based on a daily or monthly frequency; and whether they include each macroeconomic news indicator separately or jointly in a multivariate set-up. Accordingly, the results lend support to an identification strategy that orders oil before the variables measuring economic activity and inflation.

Here, we replicate this exercise for the euro area sample considered in the current paper (focusing on the impact of news regarding inflation and industrial production, which are the two macroeconomic variables included in the VAR model). An assessment of whether the evidence in support of predetermined oil prices carries over to this sample appears warranted especially in view of the more recent time period it covers compared to Kilian and Vega (2011). As documented by Fratzscher et al. (2014), the behaviour of oil prices has undergone a marked shift in the early 2000s, becoming more sensitive to changes in other asset prices since then. Our sample, which starts in 2004, falls entirely into the ‘higher-sensitivity’ period, whereas the Kilian-Vega sample for the oil-price regressions stretches from 1983 to 2008, thus predominantly consisting of observations from the ‘low-sensitivity’ period.²¹

²⁰Formally, the suite of regressions in Kilian and Vega (2011) is used to test *one* hypothesis, namely that macroeconomic news have a non-zero effect on energy prices, using *many* different specifications. This type of specification search runs counter to standard notions of statistical inference (see Leamer (1978)), and in particular may lead to an over-rejection of the null hypothesis of no effect.

²¹Like Kilian and Vega (2011), Fratzscher et al. (2014) do not find significant effects for most macroeconomic news indicators but this does not provide sufficient support for the assumption of predetermined oil prices in

The results of this exercise speak against ordering oil prices before the macroeconomic variables in our analysis. To see this, table 1 presents OLS estimates of the equation:

$$R_t = \alpha + \beta S_t^p + \gamma S_t^i + \varepsilon_t$$

where R_t is the log-change in the price of oil from the end of day $t - 1$ to the end of day t ; S_t^p and S_t^i , respectively, are news released on day t regarding euro area industrial production and headline HICP inflation (*i.e.* the two macro variables ordered before the oil prices in the baseline specification introduced in section 2.3); β and γ are the slope coefficients; α is an intercept and ε_t an error term. Following the approach in Kilian and Vega (2011), the news variables are calculated as the difference between the data release and the median expectation from Bloomberg surveys for the respective variable and normalised by the sample standard deviation of that difference to facilitate the interpretation of coefficients. The data range from February 2004 to mid-September 2016 (for descriptive statistics see table 2).

Table 1: Regression estimates (oil price changes as dependent variable)

	in USD	in EUR
Industrial production	-0.013 (0.089)	-0.073 (0.091)
HICP inflation	0.142* (0.076)	0.145* (0.084)
Constant	0.019 (0.112)	-0.065 (0.122)
Observations	292	292
R^2	0.006	0.006

* p<0.10 ** p<0.05 and *** p<0.01. Heteroskedasticity robust standard errors in parentheses.

While surprises in industrial production do not affect oil prices, the estimates point to a positive coefficient on HICP surprises that is significant at a 10% level (against a two-sided alternative), independent of whether oil prices are expressed in USD and EUR terms.²² Notwith-

the context of our analysis: first, Fratzscher et al. (2014) include the news indicators alongside a range of other financial asset prices that may immediately reflect any macroeconomic news and may hence conceal significant effects due to multi-collinearity; our model, by contrast, does not include some of these financial asset prices (in particular equities which are likely to be particularly sensitive to macroeconomic news) and may hence point to different conclusions. Second, Fratzscher et al. (2014) include a wider range of macroeconomic news variables beyond those relevant to our model which may reinforce multi-collinearity issues.

²²Testing H_0 against a *one-sided alternative*, as in Kilian and Vega (2011), would render γ significant at a 5% level, while leaving the conclusions unaltered regarding β . Unlike Kilian and Vega (2011), we do not compute ‘data-mining robust standard errors’ since the regressions are based on only one specification to test the significance

standing the relatively limited quantitative effect and an overall low explanatory power of the regression, the significant effect of HICP surprises indicates that it is more prudent to not treat oil prices as predetermined in the MS-VAR introduced in section 2.3.

Table 2: Summary statistics

Variable	Mean	Std. Dev.	Min.	Max.
Oil price in USD	0.018	1.888	-7.572	8.017
Oil price in EUR	-0.06	2.045	-11.507	9.164
Industrial production	-0.105	1	-4.482	4.98
HICP inflation	-0.019	1	-2.768	2.768

Notes: Oil price variables defined as 100 times log change in the oil price level from end of day $t - 1$ to end of day t . Industrial production and HICP inflation defined as difference between release and ex ante expectations normalised by standard deviation; see text for additional detail.

of the two macro variables relevant for the MS-VAR. This is in contrast to Kilian and Vega (2011) who estimate multiple specifications to test the hypothesis that any element of a set of 30 macroeconomic news variables is significant; it is this repeated application of the same statistical test that motivates their use of data-mining robust standard errors.

A.2 Additional figures

Figure 5: Impulse responses to oil price shock (one standard deviation shock), slashed blue: normal regime, solid red: adverse regime

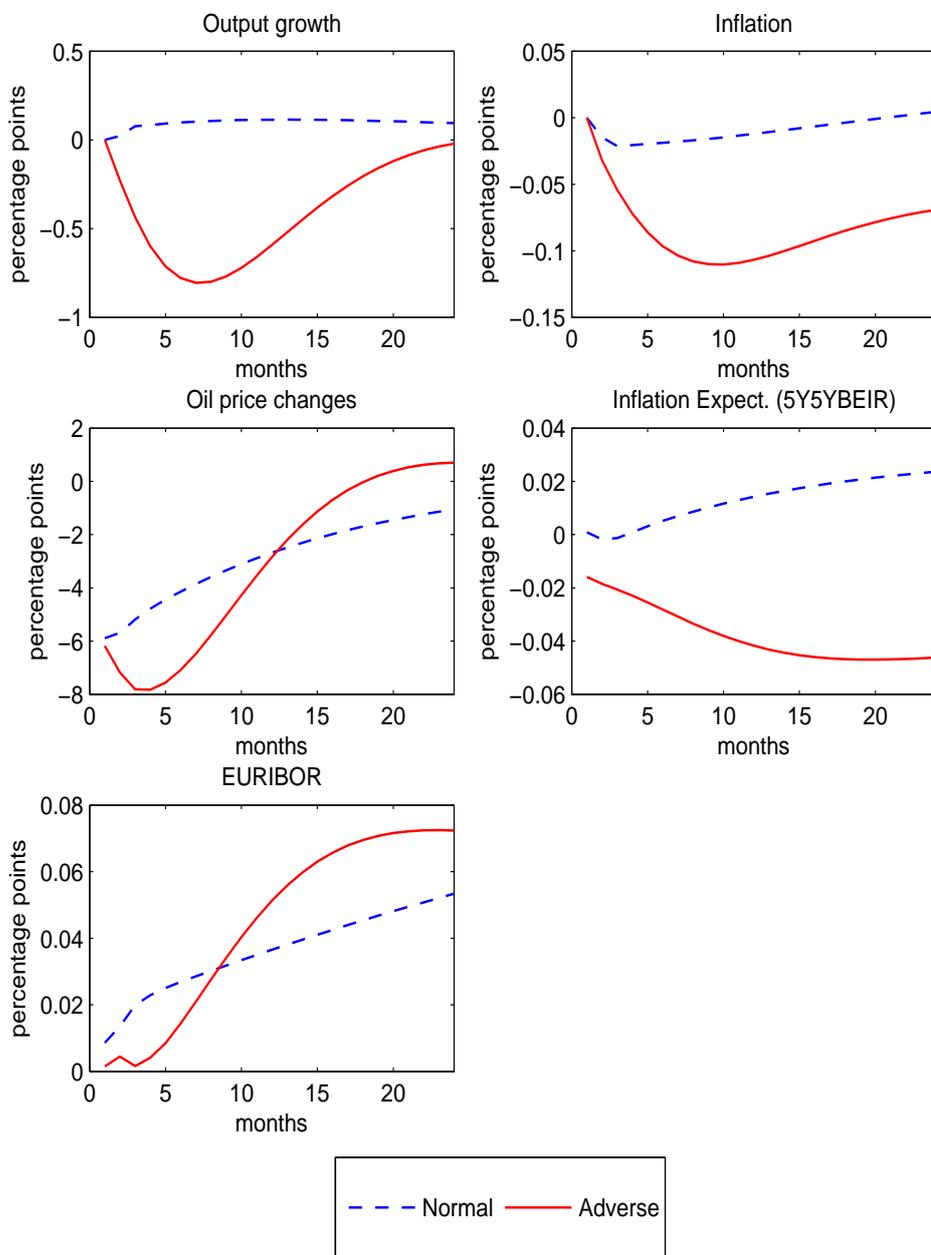
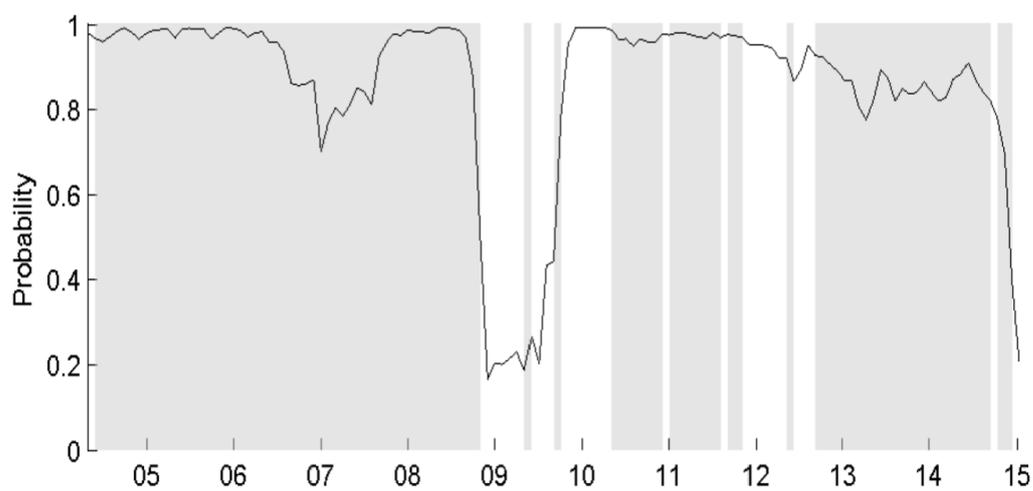


Figure 6: Time-varying probability for being in a normal regime (grey-shaded area) and conditional probability of staying in that regime (black line) – oil price in euro



Note: The ticks on the x-axis refer to the beginning of the respective year (with “05” indicating the year 2005 etc.)

Figure 7: Impulse responses to oil price shock in model with nominal negotiated wages (one standard deviation shock), slashed blue: normal regime, solid red: adverse regime

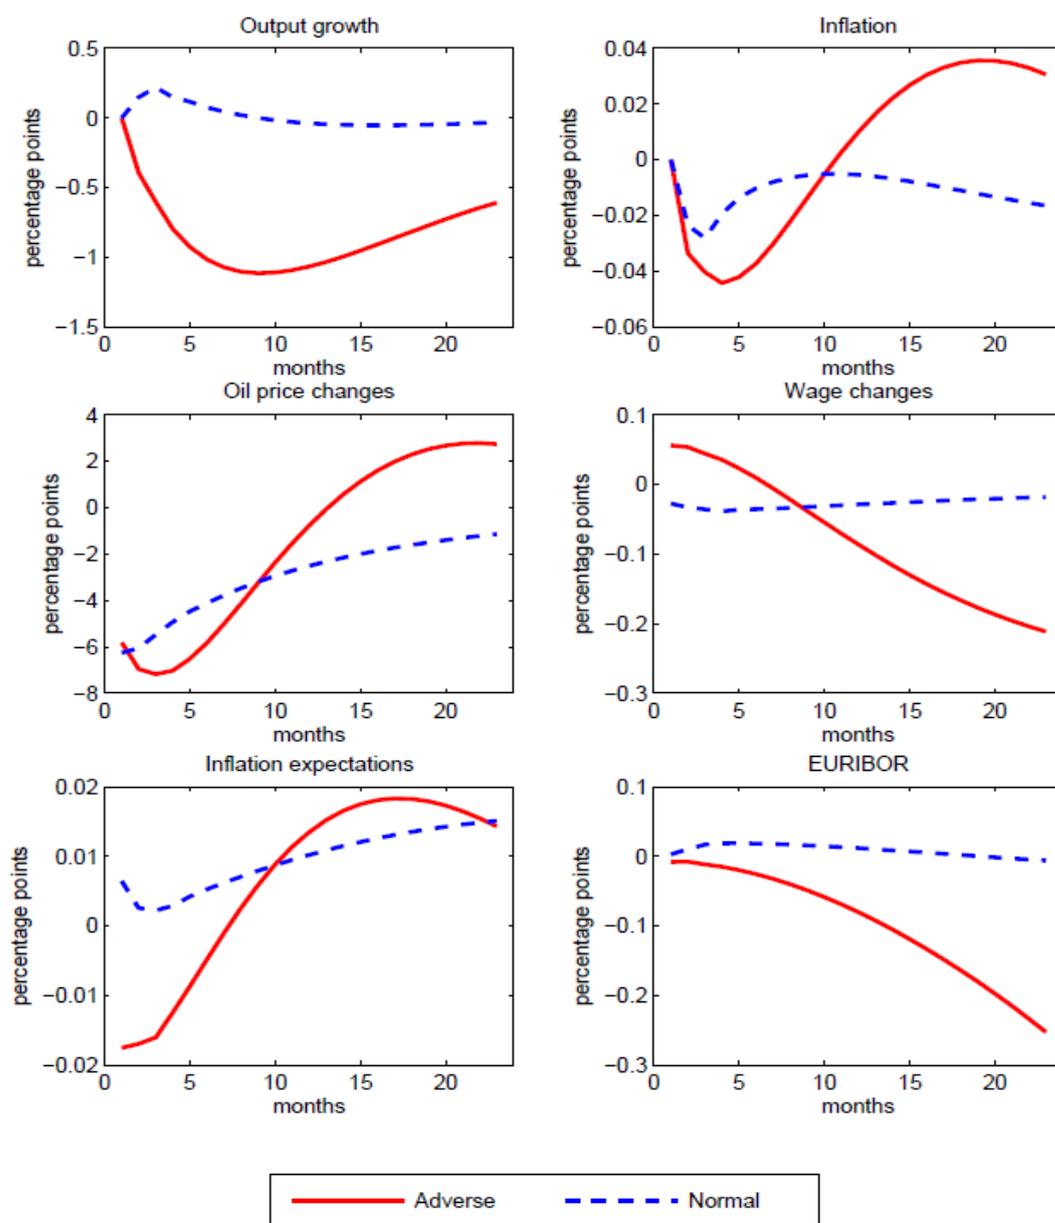


Figure 8: Impulse responses to oil price shock in model with long-term real interest rate (one standard deviation shock), slashed blue: normal regime, solid red: adverse regime

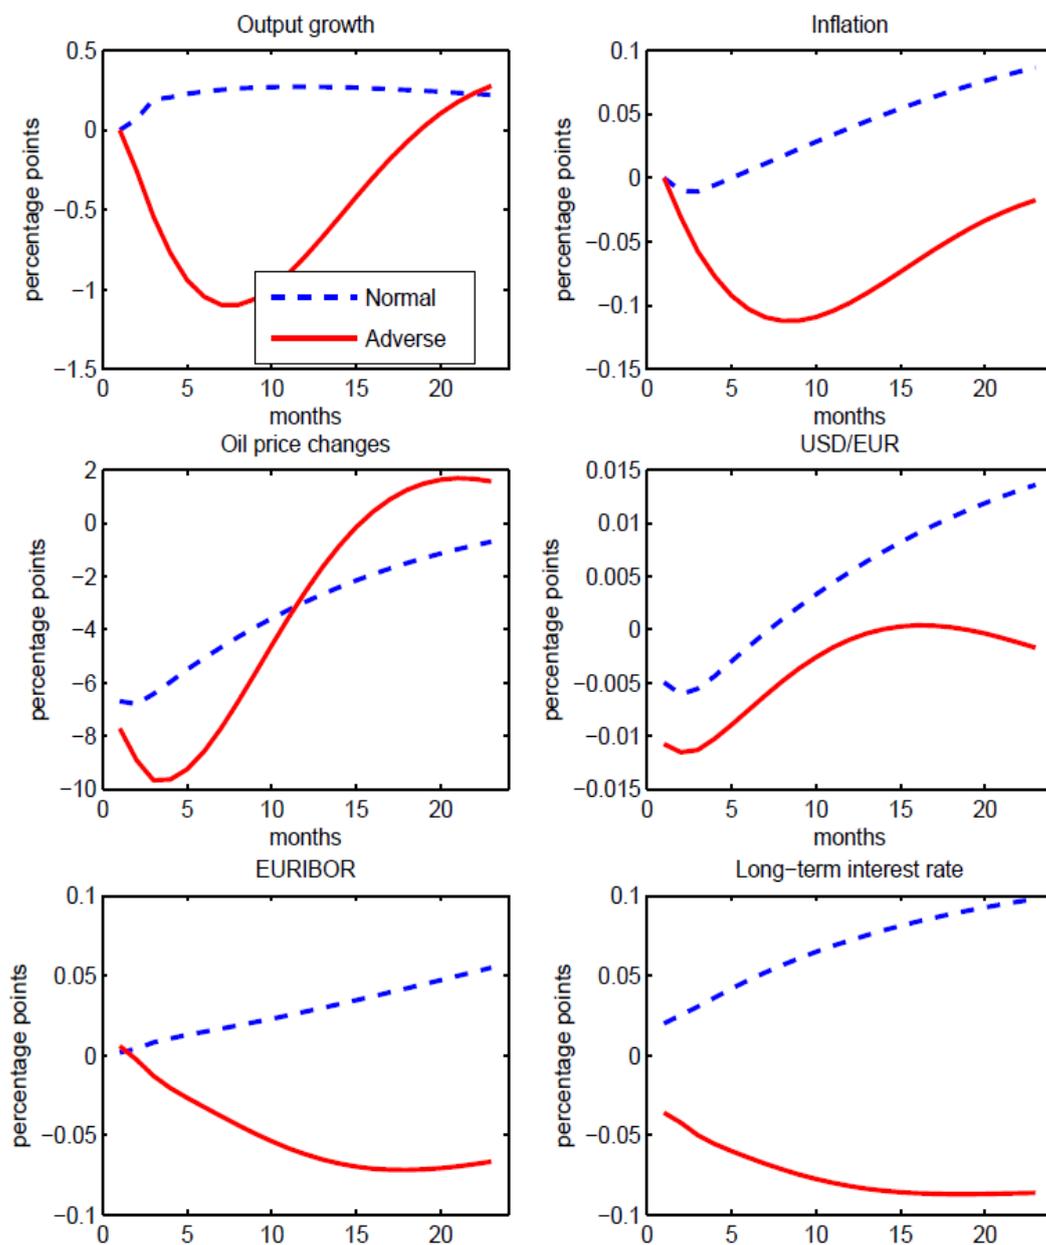


Figure 9: Impulse responses to oil price shock for 1974-2015 sample, slashed blue: normal regime, solid red: adverse regime

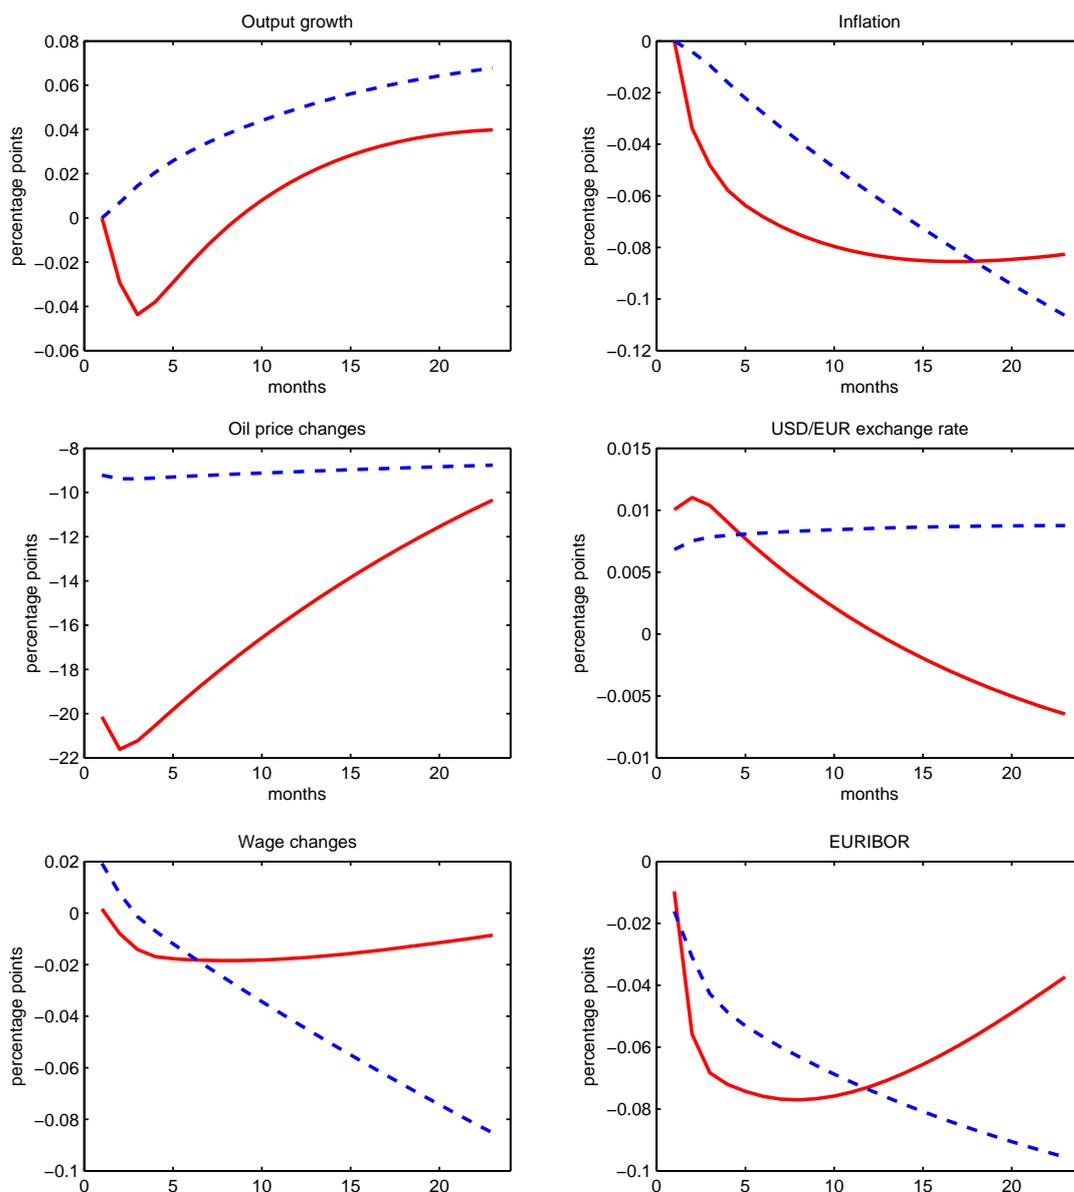
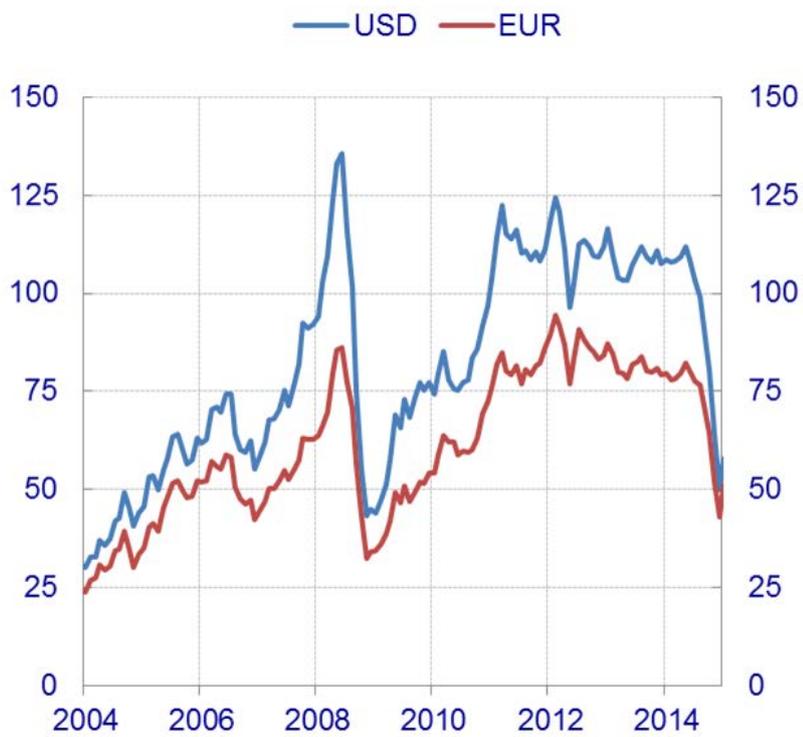


Figure 10: Evolution of Brent crude oil prices (in different currencies)



Source: *Bloomberg*

Acknowledgements

We thank Lutz Kilian, Hans-Joachim Klöckers, Wolfgang Lemke, Simone Manganelli, Roberto Motto, Ulrich Müller, Massimo Rostagno, Chris Sims, Robert Tetlow, Robert Vigfusson, Daniel Waggoner, Mark Watson, Tao Zha, and conference participants at the International Association for Applied Econometrics Conference 2016, the Oxmetrics Conference for Macroeconomics and Finance 2016, the 10th International Conference on Computational and Financial Econometrics, and the 21st Annual International Conference on Macroeconomic Analysis and International Finance, as well as seminar participants at Princeton University, the European Central Bank and the Federal Reserve Board for useful comments. We also thank Nikolaj Broberg, Mattia Colombo, Nick Ligthart and Kevin Starnes for excellent research assistance. The views expressed here are those of the authors and do not necessarily reflect those of the European Central Bank or the Federal Reserve System or the staff of either of these institutions.

Fédéric Holm-Hadulla

European Central Bank, Frankfurt am Main, Germany; email: federic.holm-hadulla@ecb.europa.eu

Kirstin Hubrich

Federal Reserve Board, Washington, United States; email: kirstin.hubrich@frb.gov

© European Central Bank, 2017

Postal address 60640 Frankfurt am Main, Germany
Telephone +49 69 1344 0
Website www.ecb.europa.eu

All rights reserved. Any reproduction, publication and reprint in the form of a different publication, whether printed or produced electronically, in whole or in part, is permitted only with the explicit written authorisation of the ECB or the authors.

This paper can be downloaded without charge from www.ecb.europa.eu, from the [Social Science Research Network electronic library](#) or from [RePEc: Research Papers in Economics](#). Information on all of the papers published in the ECB Working Paper Series can be found on the [ECB's website](#).

ISSN	1725-2806 (pdf)	DOI	10.2866/885147 (pdf)
ISBN	978-92-899-3053-6 (pdf)	EU catalogue No	QB-AR-17-131-EN-N (pdf)