# Oil Prices Uncertainty, Endogenous Regime Switching, and Inflation Anchoring*

Yoosoon Chang †
Indiana University

Ana María Herrera ‡
University of Kentucky

Elena Pesavento §
Emory University

February 3, 2023

**Abstract**

Using a novel approach to model regime switching with dynamic feedback and interactions, we extract latent mean and volatility factors in oil price changes. We illustrate how the volatility factor constitutes a useful measure of oil market risk (or oil price uncertainty) for policy makers and analysts as it captures uncertainty not reflected in other economic/financial uncertainty measures. Then, in the context of a VAR, we investigate the role of oil price uncertainty in driving inflation expectations and inflation anchoring. We show that shocks to the mean factor lead to higher expected inflation and inflation disagreement among professional forecasters and households. In contrast, shocks to the volatility factor act as aggregate demand shocks in that they result in lower expected inflation, yet they do increase disagreement about future inflation among professional forecasters and, especially, among households. We also provide econometric evidence suggesting the proposed endogenous volatility switching model can outperform other regime switching models.

JEL Classification: C13, C32, E32, Q35.
Key words: oil price volatility, endogenous regime switching, expected inflation, inflation anchoring

*We are grateful to two anonymous referees, the Associate Editor, and participants in the 2019 Workshop on Energy Economics and the 2019 Midwest Econometrics Group meetings for very useful comments. Our special thanks go to Sangmyung Ha and Diego Rojas Baez for their superb research assistance and to Joon Y. Park and Shi Qiu for many helpful discussions.

†Indiana University, Email: yoosoon@indiana.edu.
‡University of Kentucky, Email: amherrera@uky.edu.
§Emory University, Email: epesave@emory.edu.

# 1 Introduction

Historically, crude oil prices have exhibited periods of increased volatility and periods of relative calm. Spikes in oil prices have been observed in the heels of political unrest in the Middle East, terrorist attacks on Saudi Arabia's oil facilities, and have often coincided with economic downturns in the U.S. economic activity such as the global shutdown that ensued the onset of the Covid-19 pandemic. Over the years, concerns regarding heightened volatility and unexpected increases in oil prices –as well as the burdens they may pose on households, businesses and investors– have been reflected in policy statements by different chairmen of the Federal Reserve System and other central banks, demands from senators to the U.S. Commodity Futures Trading Commission, and an extensive literature on the economic impact of oil price shocks.

How do we best model the transition from periods of calm to periods of turmoil in the oil market? Do oil price fluctuations and oil price uncertainty affect inflation and the inflation anchor? Work by Coibion and Gorodnichenko (2015) provides empirical evidence that oil price increases between 2009 and 2011 played a key role in the rise of household inflation expectations during the period of moderate inflation. But could unexpected increases in the rate of growth of oil prices or heightened oil price uncertainty result in starker disagreement among forecasters or household inflation expectations?

The objective of this paper is to address these questions in the context of structural vector autoregressive (SVAR) model while at the same time deriving a measure of oil market risk that could be of use to policy makers and analysts. We start by estimating a Markov switching model with dynamic feedback and interactions to inquire into the process that drives transitions from periods of calm to periods of turmoil in oil markets. Our inquiry into these transitions is prompted by the historical behavior of the West Texas Intermediate (WTI) and its percentage change depicted in Figure 1. Crude oil prices, as other financial asset prices, often experience abrupt fluctuations. For instance, as Figure 1 illustrates, the price of the WTI rose steadily from a \$19.46 per barrel in November 2001 to a peak of \$139.96 in June 2008, and then it plunged to \$41.96 in January 2009. As for the volatility of the WTI, it was moderate during the mid and late 1990s, increased considerably in the early 2000s and skyrocketed during the Covid-19 crisis.

To study oil price dynamics, we estimate an endogenous regime switching model where we allow the mean and volatility of the growth rate of the WTI to transition in an unsynchronized manner between two (high/low) states of the mean and two (high/low) states of the volatility, respectively. Hereafter, we will refer to this model as the unsynchronized mean-volatility switching model. This modeling strategy allows the mean and volatility regimes to be determined by two correlated latent factors, one driving the mean regimes and the other driving the volatility regimes, while it also accounts for possible feedback from past innovations in the growth of the WTI to both mean and volatility regime factors.

We evaluate the out-of-sample performance of the endogenous regime switching model relative to that of two close competitors: the standard regime switching model where shifts are driven by an exogenous Markov chain and the switching model with time-varying transition probabilities (TVTP) proposed by Diebold et al. (1994) where we use two alternative predetermined variables (lagged money supply, lagged inflation, and the lagged change in the WTI price.).[1]

Several insights regarding the behavior of the WTI prices are derived from our analysis.[2] First, the endogenous unsynchronized mean-volatility switching model does a better job at capturing the evolution of oil prices than an exogenous switching model. Indeed, there is a great degree of variability in the transition probabilities over time, especially around the Persian Gulf War, the Great Recession, and the onset of the Covid-19 pandemic. In addition, the endogenous volatility switching model can outperform the alternatives in forecasting out of sample.

A result of interest for policy makers, analysts, and investors is that, when the current regime is known, allowing the regime switching to evolve in an endogenous manner is very informative regarding the likelihood of whether the oil price change next period will remain in the same regime. When we examine four periods of low mean and high volatility, we find that the probability of staying in this regime declined quickly for all episodes, but the Great Recession. Thus, not all episodes should convey the same degree of "fear" to economic agents. Clearly, the Covid-19 pandemic resulted in extremely high risk levels not seen in the previous half century, as reflected in the evolution of the extracted volatility regime factor.

Could this extracted volatility regime factor be used to gauge risk in periods where alternative measures of oil market volatility are not readily available? After all, the Chicago Board Options Exchange did not publish the implied oil volatility index, OVX, before 2007. To answer this question, we compare the extracted volatility regime factor with the OVX and with commonly-used measures of economic uncertainty. We illustrate how the extracted

[1] The choice of variables is based on LASSO estimation results, available from the authors upon request, which suggest these two variables have the most explanatory power for movements in the mean and volatility of the WTI price. In this way, we try to stack the odds against our specification and towards the TVTP model.

[2] Estimation results available from the authors upon request, show similar results for Brent.

volatility factor could constitute a measure of oil market "risk" and, thus, provide analysts and investors with a tool to gauge volatility in the crude oil market.

Using a structural vector autoregressive model, we then show how the extracted regime factors can aid in understanding the response of inflation, inflation expectations and inflation disagreement to increased uncertainty in oil markets. On the one hand, we find that unexpected increases in the mean factor (i.e., faster increases in oil prices ) lead to higher inflation and short-run and long-run expected inflation, as well as higher disagreement among both professional forecasters and households. The short-run response of household inflation expectations and disagreement exceeds that of professional forecasters by an order of magnitude. This result is consistent with the findings of Coibion and Gorodnichenko (2015) who suggest that oil prices are more salient for households. On the other hand, we find that increases in the 'risk factor' lead to lower inflation as well as short-run and long-run inflation expectation; this indicates that, as other uncertainty shocks, heightened oil market uncertainty acts as a negative demand shock. Nevertheless, increases in oil price uncertainty raise inflation disagreement and, thus, can pose a threat to inflation anchoring.

Our work is closely related to two strands of literature that employ Markov switching (MS) models in the analysis of oil price fluctuations. Researchers have relied on MS models to forecast volatility of crude oil prices and compare the predictive ability of MS-GARCH and MS long memory models to selected GARCH competitors (see e.g., Di Sanzo 2018, Herrera, Hu and Pastor 2018 and references therein). This strand of literature has found evidence that MS models tend to have superior predictive ability, especially during periods of turmoil. Our paper differs in that the forecasting literature mainly relies on MS models where the transition probabilities evolve in an exogenous manner. In contrast, in the endogenous MS model employed in this paper, time-variation in the transition probabilities stems from the feedback mechanism and the dynamic interactions between the mean and volatility regime factors.[3]

Our work is also related to studies that employ endogenous MS models to investigate the relationship between oil price fluctuations and the macroeconomy.[4] For instance, Bjørnland et al. (2018) develop a Markov switching rational expectations New Keynesian model to study the role of oil price shocks in accounting for variability in the U.S. economy, as well as

[3] A brief out-of-sample forecasting evaluation shows that our volatility switching model with endogenous feedback produces a smaller RMSE than the alternative models with time-varying transition probabilities for five- and ten-year rolling window forecasting schemes.

[4] Endogenous MS models have also been used to study the transmission of financial, monetary and fiscal policy shocks. See, for instance, Davig and Leeper (2006), Benigno, et al. (2020), Hubrich and Waggoner (2021), and Chang, Kwak and Qui (2021).

the contribution of a more hawkish monetary policy regime to the decline in macroeconomic volatility. Along similar lines and using data for the euro area, Holm-Hadulla and Hubrich (2017) identify two different regimes in the response of economic activity and inflation to oil price shocks: an adverse regime, where oil price shocks result in significant and sustained changes in economic activity and inflation, and a normal regime where the response of these variables is smaller and shorter-lived.

Our paper differs from the above strand of literature in several aspects. First, we do not build a structural model of the interaction between the economy and the oil market. Hence, we do not identify a-priori the possible sources of changes in regime (e.g., monetary policy states or past inflation rates), nor do we explicitly model the interactions between oil prices and the macroeconomy. Instead, we estimate a reduced form model where we endogenize the transition probabilities by allowing them to depend only on the past behavior of the crude oil prices. Then, we use the extracted latent regime factors in a small-scale SVAR model to illustrate how our measure of oil market uncertainty can be useful to policy makers and analysts in understanding the response of inflation, inflation expectations and disagreement to increased oil market uncertainty. A disadvantage of our approach is that we cannot provide a structural interpretation of the source of oil price fluctuations. However, our approach allows us to be agnostic regarding the variables that might drive the transition between regimes. Furthermore, estimation is relatively straight-forward – we use a modified version of the Kalman filter and rely on maximum likelihood estimation instead of Bayesian methods – and it allows us to derive a measure of oil market "risk".

Finally, our paper builds on a broad line of literature exploring the drivers of inflation expectations (see Coibion and Gorodnichanko, 2015 and references therein). It is more closely related Binder (2018) and Kilian and Zhou (2022) who explore the effect of gasoline price changes on inflation expectations. Our approach differs from these studies in that we do not estimate a structural model relating gasoline prices changes and inflation expectations. Instead, we focus on the effect of the mean and volatility regime factors on inflation, inflation expectations and, especially, disagreement.

This paper is organized as follows. Section 2 describes the endogenous regime switching models and the data. Section 3 describes the estimation results for a one factor (volatility switching) and a two-factor (volatility-mean switching) models, discusses how these estimates can be used to evaluate the likelihood to remain in a high volatility or a low mean-high volatility regime, and evaluate the relative out-of-sample performance of the models. Section 4 illustrates how the extracted latent volatility regime factor may constitute a measure of oil market risk, and inquire into the effect of the increased oil price uncertainty on inflation,

inflation expectations, and disagreement. Section 5 concludes.

# 2 The Behavior of Crude Oil Prices: Endogenous versus Exogenous Regime Switching

An advantage of using a Markov switching model for studying the evolution of oil prices is that it is well-suited to investigate time series processes characterized by periods of calm and turmoil. Indeed, several studies suggest such models do a good job at tracking and forecasting the evolution of crude oil prices.[5] While conventional regime switching models are well-understood and econometric toolboxes are readily available for their estimation, conventional models do not allow feedback from past oil price changes into the transition probabilities. As such, they do not provide researchers with an effective tool to infer the likelihood of transitioning to a different state when the current state is known. In this section we briefly describe the endogenous Markov-switching models, describe the data and the estimation strategy.

## 2.1 Endogenous versus Exogenous Regime Switching Models

Consider an endogenous regime switching model similar to Chang, Choi and Park (2017) – hereafter CCP –, which differs from their specification in that the mean and volatility of crude oil price changes are allowed to switch in an unsynchronized manner. Four regimes for the log difference in oil prices $y_t$ are possible:[6] (1) low-volatility, low-mean; (2) low-volatility, high-mean; (3) high-volatility, low-mean; and (4) high-volatility, high-mean. Such a model is given by

$$y_t - \mu\left(s_{m,t}\right) = \sum_{k=1}^{p} \gamma_k \left(y_{t-k} - \mu\left(s_{m,t-k}\right)\right) + \sigma\left(s_{v,t}\right) u_t \tag{1}$$

The parameters $\mu_{s_{m,t}}$ and $\sigma_{s_{v,t}}$ denote the time-varying conditional mean and volatility of the oil price changes that depend on two distinct, but correlated state processes $s_{m,t}$ and $s_{v,t}$. The process $s_{m,t}$ specifies the binary state of the mean, with $s_{m,t} = 0$ and 1 respectively representing *low* and *high* mean states. Similarly, the binary volatility state process $s_{v,t}$ specifies *low* and *high* volatility states with $s_{v,t} = 0$ and 1.[7]

[5] See Herrera, Hu and Pastor (2018) for results of Markov switching tests indicating such a model is appropriate in this context.

[6] Throughout the text we use the terms oil price change and growth in oil prices to describe $y_t$.

[7] For regime identification, we assume $\mu(0) < \mu(1)$ and $\sigma(0) < \sigma(1)$.

We specify the state processes $s_{i,t}$ for $i = m, v$ as

$$s_{i,t} = 1\left\{w_{i,t} \geq \tau_i\right\}$$

with a latent regime factor $w_{i,t}$ and a threshold $\tau_i$. The mean (volatility) regime factor $w_{m,t}$ ($w_{v,t}$) determines the switch between states of low and high mean (volatility) according to whether it is below or above the threshold $\tau_m$ ($\tau_v$). We let $\boldsymbol{w_t}= (w_{m,t}, w_{v,t})'$ and jointly consider the dynamics of the two latent factors by assuming they follow a first-order stationary bivariate autoregressive process

$$\boldsymbol{w_t} = \boldsymbol{A}\boldsymbol{w_{t-1}} + \boldsymbol{v_t}$$

where

$$\boldsymbol{A} = \begin{pmatrix} a_{mm} & a_{mv} \\ a_{vm} & a_{vv} \end{pmatrix},$$

the modulus of all eigenvalues of $\boldsymbol{A}$ is less than unity, the innovations $\boldsymbol{v_t}= (v_{m,t}, v_{v,t})'$ are independent and identically distributed over time and correlated with the previous oil price change innovation $u_{t-1}$. Specifically, we assume $(u_{t-1}, v_t')' \sim i.i.d.\mathbb{N}(0, \mathbf{P})$ with a correlation matrix

$$\boldsymbol{P} = \begin{pmatrix} 1 & \boldsymbol{\rho_{vu}'} \\ \boldsymbol{\rho_{vu}} & \boldsymbol{P_{vv}} \end{pmatrix} = \begin{pmatrix} 1 & & \\ \rho_{v_m,u} & 1 & \\ \rho_{v_v,u} & \rho_{v_m,v_v} & 1 \end{pmatrix} \tag{2}$$

where variances are normalized for identification.

Note that we could restrict the switch in the volatility and the mean to be synchronized, in which case the model is given by[8]

$$y_t - \mu\left(s_t\right) = \sum_{k=1}^{p} \gamma_k \left(y_{t-k} - \mu\left(s_{t-k}\right)\right) + \sigma(s_t)u_t. \tag{3}$$

Then, a single state, $s_t$, and latent factor, $w_t$ govern switches in both $\mu$ and $\sigma$, where $w_t = \alpha w_{t-1} + v_t, |\alpha| < 1$, and $\begin{pmatrix} u_{t-1} \\ v_t \end{pmatrix} \sim i.i.d.\,\mathbb{N}\left(0, \begin{pmatrix} 1 & \rho_{v,u} \\ \rho_{v,u} & 1 \end{pmatrix}\right)$. We refer to that model as the synchronized switching model. We could further assume that only the volatility switches between regimes, in which case the model would simplify to

$$y_t = \mu + \sum_{k=1}^{p} \gamma_k \left(y_{t-k} - \mu\right) + \sigma\left(s_t\right) u_t \tag{4}$$

[8]This restricted model is a generalized version of the Markov switching model proposed by Hamilton (1989, 2010), where the regime switches are endogenized via the covariance between $u_{t-1}$ and $v_t$, $\rho_{u,v}$.

where the latent factor, and innovations are specified as above.[9] We refer to this model as the volatility switching model. When $\rho_{v,u} = 0$, the model in (4) is equivalent to the exogenous Markov switching model for volatility.

However, a-priori, there is no reason to believe that only the volatility switches or that switches in the mean and volatility are synchronized. For instance, after the oil price collapse of 1986, the rate of growth of oil prices was below the historical mean while volatility was high. Volatility also skyrocketed during the Covid-19 pandemic, whereas oil prices plummeted. In contrast, at the beginning of the Persian Gulf war both the mean and volatility of oil prices increased.

To get a better grasp of what this generalization of the CCP model entails for the dynamics of crude oil price changes, let us examine the role of the different parameters. In the unsynchronized model (1), the evolution of the bivariate latent regime factor $\boldsymbol{w_t} = (w_{m,t}, w_{v,t})'$ is driven by the innovations $v_{m,t}$ and $v_{v,t}$ collected in the vector $\boldsymbol{v_t} = (v_{m,t}, v_{v,t})'$, and by the dynamic interaction between the two factors captured by the autoregressive coefficient matrix $\boldsymbol{A}$ and their contemporaneous correlations given by the correlation matrix $\boldsymbol{P}$. In particular, if $a_{mv} \neq 0$, then the volatility regime factor helps to predict the mean regime factor. Conversely, if $a_{vm} \neq 0$, then the mean regime factor helps to predict the volatility regime factor. When $\boldsymbol{A}$ is diagonal, larger values of $a_{mm}$ and $a_{vv}$ , respectively, would indicate higher persistence in the mean and volatility regime factors. The correlation parameters $\rho_{u,v_m}$ and $\rho_{u,v_v}$ capture feedback from past oil price change innovations, $u_{t-1}$, to the mean and volatility factors. Endogenous feedback may occur through two different channels. Shocks to crude oil price changes at time $t$ would affect the regime switching in the mean at time $t+1$ if $\rho_{u,v_m} \neq 0$. Furthermore, shocks to rate of growth of oil prices at time $t$ could also affect the regime switching in the volatility at $t+1$ if $\rho_{u,v_v} \neq 0$. Contemporaneous correlation between the factor innovations $v_{m,t}$ and $v_{v,t}$ net of the contribution of $u_{t-1}$ is given by $\rho_{v_m v_v \cdot u} = \rho_{v_m,v_v} - \rho_{v_m,u}\rho_{v_v,u}$. Thus, we may examine a null hypothesis of no contemporaneous or dynamic interaction between the mean and volatility regime factors by testing whether $\rho_{v_m v_v \cdot u} = 0$ or $a_{vm} = a_{mv} = 0$. Clearly, these feedback channels are absent in the conventional Markov switching model.

The endogenous Markov switching model estimated in this paper is similar to those proposed by Diebold et al. (1994) and Filardo (1994) in that it allows the transition probabilities to be time varying. Both papers model the transition probabilities as functions of strictly exogenous explanatory variables and lagged values of the dependent variable. We also assume

[9] Given the assumption that $|\alpha| < 1$, the latent regime factor $w_t$ is stationary.

the probabilities are time varying and may depend on lagged values of the shocks to the dependent variable. However, we do not take a stand on the functional form or exogenous and predetermined variables that derive the dynamics of the transition probabilities. Here, the transition probabilities are a function of unexpected shocks to oil prices; yet we do not explicitly specify variables or events that drive these shocks. Moreover, large shock-inducing events (e.g., the Persian Gulf war) may be understood as innovations to changes in the oil price that are not captured by the time-varying mean and volatility, but that are captured by the oil price surprises, $u_t$.

## 2.2 Data and Empirical Methodology

To study the behavior of crude oil prices we use the monthly spot price of the West Texas Intermediate (WTI) spanning the period between January 1986 and January 2021 obtained from the Energy Information Agency.[10] We compute oil price changes (i.e., the rate of growth) as the first difference of the natural log of the spot price. We focus on the WTI price as it is the reference price for buyers and sellers of crude oil in the U.S., it is produced in the U.S. and – until late 2015 when the export ban on U.S. crude oil was lifted – was only sold in the U.S. Nevertheless, we note that estimation results for the endogenous Markov switching models are robust to using Brent crude oil price changes.[11]

Crude oil prices, as other financial assets, often experience abrupt changes in behavior. As Figure 1 illustrates, the price of the WTI, denoted by the blue line, exhibited an upward trend during the 1990s and most of the 2000s, reaching a peak of $139.96 per barrel in June 2008. Then, it plunged during the Great Recession to a trough of $41.96 per barrel in January 2009, slowly recovering during the first half of the 2010s until it collapsed in July of 2014. A striking decline was also observed in the wake of the global shutdown induced by the Covid-19 pandemic.[12]

FIGURE 1 HERE

Percentage changes in the WTI (Figure 1) suggest periods of political and economic turmoil might be associated with higher volatility. See, for instance, the period surrounding the Persian Gulf War (August 1990 to February 1991), the global financial crisis in the late

[10] Oil prices were controlled in the U.S. prior to 1973. Studies that use prices for the 1970s –or earlier periods– rely on splicing data for the WTI and the producer price index or the refiners acquisition cost. Moreover, the WTI was not adopted as a benchmark price until the 1980s.

[11] For the sake of brevity, we relegate the results for Brent oil price to the online appendix.

[12] Figure A.1 of the online appendix reveals a similar behavior for Brent.

2000s, and at the onset of the Covid-19 pandemic. Yet, periods of high volatility do not always coincide with periods of high mean. For instance, both oil prices and their volatility were high during the period following the outbreak of the Persian Gulf war; however, the 2008 financial crisis was characterized by low mean and high volatility.[13]

Regime switching models can be estimated via maximum likelihood. We refer the reader to CCP for a detailed description of the estimation procedure for regime switching models with a single latent factor such as (3) and (4). For the estimation of our general unsynchronized mean and volatility switching model (1), we use the algorithm and filter developed for regime-switching models with multiple latent regime factors by Chang, Kwak, Park and Qiu (2021). As discussed in CCP, when the state is correlated with the observed process, namely the percentage change in crude oil prices, a modified filter is needed to explicitly account for this endogenous feedback channel that results in time-varying transition probabilities. To estimate our model (1) with two regime factors, we use the Chang-Kwak-Park-Qiu filter, which extends the CCP filter to a multivariate setup. As in the standard Kalman filter, our modified filter involves the usual prediction and updating steps, but they are carried out with time-varying transition probabilities.

To illustrate how the endogenous feedback gives rise to time-varying transition probabilities, first we net out the effect of the past oil price change innovation $u_{t-1}$ from the regime factor innovations to obtain the orthogonal regime factor shock $\mathrm{v}_{\boldsymbol{t}}$ - $\boldsymbol{\rho_{vu}}\mathrm{u}_{t-1} \sim \mathbb{N}(0, \boldsymbol{P_{vv\cdot u}})$. Then, we use the conditional distribution of $v_t$ given $u_{t-1}$ to compute the transition probability. For example, the transition probability from the low mean-high volatility regime to the low mean-high volatility regime (LH-to-LH), or staying in the LH regime, is

$$\mathbb{P}\left\{\boldsymbol{s_t} = (0,1)' | \boldsymbol{s_{t-1}} = (0,1)', \mathcal{F}_{t-1}\right\} = \Phi(\boldsymbol{\tau})^{-1} \int_{-\infty}^{\tau} \Phi_{v|u}\left(\boldsymbol{\tau} - \boldsymbol{\rho_{vu}} u_{t-1} - \boldsymbol{A}\boldsymbol{w_{t-1}}\right) \phi(\boldsymbol{w_{t-1}}) d\boldsymbol{w_{t-1}}$$

where $s_t = (s_{m,t}, s_{v,t})'$, $\Phi_{v|u}$ is the conditional distribution the regime factor innovation $v_t$ given the past innovation to oil price changes $u_{t-1}$, $\boldsymbol{\tau}$=$(\tau_m, \tau_v)'$ and $\mathcal{F}_{t-1}$ is the information set available at time $t-1$ given by the past oil price changes, $y_{t-1}, y_{t-2}, \ldots, y_1$. The vector $\boldsymbol{\rho_{vu}}\mathrm{u}_{t-1} = (\rho_{v_m,u} u_{t-1}, \rho_{v_v,u} u_{t-1})'$ contains the feedback effects from $u_{t-1}$ to the mean regime factor $w_{m,t}$ $(\rho_{v_m,u} u_{t-1})$, and to the volatility regime factor $w_{v,t}$ $(\rho_{v_v,u} u_{t-1})$. These feedback channels make the otherwise constant transition probabilities time-varying. The final impact of a one-unit realized oil price change innovation $u_{t-1}$ therefore depends on the sign and magnitude of its correlation with the mean and volatility regime factor innovations $v_{v,t}$ and $v_{v,t}$. Together they influence the regime determination process by effectively lowering or

[13]The same periods of high volatility are observable for Brent. See Figure A.1 of the online appendix.

raising the threshold. When the feedback channels are shut down, $\rho_{vu} = 0$, then the model reduces to the conventional Markov switching model with constant transition probabilities.

# 3 Regime Switching in the Volatility and Mean of Crude Oil Prices

While our estimates clearly identify two separate latent factors for the mean and volatility of crude oil price changes, we first build some intuition on why the behavior of crude oil price changes may be better captured by an endogenous regime switching model, by looking at the results of the volatility switching model (4). Parameter estimates with 68% confidence intervals[14] from the simple model with volatility switching only are reported in the first and second columns of Table 1. Two regimes, low and high volatility, are clearly identified: the volatility during periods of turmoil ($\sigma_1 = 0.222$) is more than thrice the volatility in periods of calm ($\sigma_0 = 0.069$) and the latent factor is very persistent ($\alpha = 0.971$). The negative and significant coefficient on $\rho_{v,u}$ is evidence of a strong leverage effect: the estimate of $\rho_{v,u} < 0$ ($-0.974$) indicates that a negative shock to the mean in period $t$ implies an increase in volatility in period $t + 1$.[15]

TABLE 1 HERE

As Figure 2 illustrates, several periods of high volatility (shaded in gray) are identified over the sample. These periods correspond to the times when the extracted latent factor exceeds the threshold, $\tau$ (red line). Three features stand out. First, regimes with high volatility are recurrent, but short-lived. We observe 29 months (out of 419 observations) in the high volatility regime, which appear to be concentrated in five different periods. Second, not surprisingly, periods of political unrest in key oil producer countries (e.g., the invasion of Kuwait) and economic contractions (e.g., the 2008 Great Recession and the 2020 Great Shutdown) constitute periods of heightened risk in the oil market. Finally, of special interest are two other episodes related to recent developments in the oil market: the increased financialization of the oil market in the early 2000s[16] and the 2014 collapse. The

[14] Confidence intervals are obtained using the stationary block bootstrap procedure by Politis and Romano (1994). We obtain percentile bootstrap confidence intervals by estimating 500 block bootstrapped samples of length 420 (404 for the Brent oil price change). The average block size is 17 for the stationary block bootstrap, which is selected by averaging the optimal block size for each vector of time series.

[15] Estimation results for Brent can be found in Table A.1 of the online appendix.

[16] The presence of financial investors has increased considerably since the early 2000s. Financial players without an interest in holding physical crude oil (e.g, pension funds, hedge funds, insurance companies) have

latent factor increases only slightly and for a brief period of time around the financialization of the oil market (but does not approach the threshold for a high volatility regime); yet it raises significantly and for a prolonged period of time after the 2014 oil price collapse. The latter clearly reflects a period of heightened volatility and increased risk.[17]

FIGURE 2 HERE

Figure 3 evidences the difference between the transition probabilities estimated from the endogenous and exogenous regime switching models. The left panel reports the probability of staying in the high volatility state and the right panel illustrates the transition probability from the low to the high volatility state. The black solid line represents the time-varying transition probability estimated from the endogenous switching model, and the red dashed line corresponds to the constant transition probability from the exogenous regime switching model. Note that, in the exogenous model, the probability to stay in the high regime remains constant at 0.86. In contrast, in the endogenous model the probability varies over time with the realized values of the oil price change, and differs significantly across various periods of high volatility. For instance, on the one hand, the probability of remaining in the high volatility regime drops to 0.39 and 0.27 during the Persian Gulf War and the Covid-19 Pandemic, respectively, suggesting both events would lead only to a temporary increase in volatility. On the other hand, the endogenous time-varying transition probabilities remained high during the Great Recession and the 2014 oil price collapse. Perhaps more striking is the difference between the transition probability from the low to the high-volatility regime in the exogenous and endogenous model. Whereas in the exogenous switching model the low to high transition probability would have remained constant at 0.021, the endogenous model reveals low to high transition probabilities that vary significantly over time and exceed 0.15 around the invasion of Kuwait, during the Great Recession, and when oil prices plunged after the second half of 2014.

Of particular interest are the transition probabilities during the Covid-19 pandemic: the low-to-high transition probability increased to about 70% at the onset of the global shutdown while the high-to-high probability declined below 30%. This behaviour suggests a period of very high uncertainty in oil markets. As we will see later, this impression is corroborated by the behavior of the oil market volatility index produced by the Chicago Board Options

since held larger positions in derivatives and futures markets. These developments have led to a heated debate regarding the role of financial speculation in driving oil price volatility.

[17]Estimation results lead us to identify almost identical periods of high volatility for Brent. See Figure A.2 of the online appendix.

Exchange (the CBOE OVX).

FIGURE 3 HERE

The first panel of Figure 3 illustrates how, once the oil price change entered a period of high volatility (denoted by the gray shaded area), the probability of remaining in the high-volatility state quickly declined for all episodes, but the Great Recession. In fact, only for the period corresponding to the Covid-19 pandemic did the endogenous switching model estimate a high-to-high probability that temporarily exceeded the estimate obtained from the exogenous model. This coincides with a period in which the WTI declined for several months (see Figure 1).[18]

To summarize, the information contained in the time-varying transition probabilities could be useful for policy analysts and investors. For instance, during periods of turmoil when volatility is high, an endogenous regime switching model could aid in assessing the risk of remaining in the high volatility state. Regardless of whether we use the WTI or Brent, the endogenous switching model provides useful information regarding oil market risks possibly associated with political unrest and economic downturns reflected already in the oil price. We will turn back to this issue in Section 4.2.

## 3.1 Unsynchronized Switching in the Mean and Volatility of Crude Oil Prices

Is the behavior of oil prices better captured by model (4) or by a model that also allows the mean to switch as in (3) or (1)? To answer this question, we first estimate the endogenous regime switching model in (1) where the switches in mean and volatility are driven by two different – yet correlated – latent factors, and then proceed to test a series of hypotheses that speak to the fit of the general model.

The third column of Table 1 reports the estimates (with 68% confidence intervals in the fourth column) from the unsynchronized model.[19]

The estimates for $\sigma_0$ and $\sigma_1$ are similar to those obtained in the simpler model with volatility switching only and also indicate that the volatility in oil price changes is more than three times as large in the high-volatility regime ($\sigma_1 = 0.219$) compared to that in the

[18] Similar results are found when using the Brent crude oil price. See Figure A.3 of the online appendix.

[19] Results for a model where the mean and the volatility switch in a synchronized manner are available from the authors upon request

low-volatility regime ($\sigma_0 = 0.059$). The fact that the estimates of $\sigma_0$ **and** $\sigma_1$ do not change much from model (4) indicates that the difference in the volatility states is not driven by the mean.

We find evidence of endogeneity both for the volatility and the mean regimes. The estimate of $\rho_{u,v_v}$, which measures the degree of endogenous feedback to the volatility regime switching, is negative and significant ($-0.917$), that of $\rho_{u,v_m}$, which accounts for endogeneity in the mean regime, is positive and significant (0.529).

FIGURE 4 HERE

Figure 4 depicts the extracted latent factors for the mean (left panel) and volatility (right panel). The mean regime factor is not very persistent. Yet, it remains below the threshold for a few months after the end of the Persian Gulf war and then again during the Great Recession, the 2014 oil price collapse, and the Great Shutdown. About 91% of the observations in the sample belong in the high-mean regime. The volatility regime factor is roughly consistent with the factor extracted from the simpler volatility switching model (see Figure 2), although the high volatility regime appears to be slightly less prevalent here (5% of the observations are in the high-volatility regime versus 7% in the model with volatility switching only).[20]

A comparison of Figure 2 and Figure 4 suggest their informational content is very similar. Given this similarity, the reader may wonder how the mean and volatility latent regime factors are related. To explore this issue, we report the 24-month rolling window correlations among the latent factors as well as the coherence (see Figure 5).[21] Three features are noticeable in this figure. First, the correlation is negative throughout the sample but exhibits a large degree of time variation. Second, although the correlation between the latent factors drops significantly during the Gulf War, it remains high throughout the Great Recession, during the 2014 oil price collapse and the Great Shutdown. Third, the coherence plot indicates that the co-movement among the two factors is accounted for, slightly more, by lower than higher frequencies, especially the frequencies corresponding to periods longer than three months. Very similar patterns emerge when we redo the analysis using Brent oil prices

[20] Figure A.5 of the online appendix illustrates the results for the Brent.

[21] The coherence between the mean and volatility regime factors, $w_{m,t}$ and $w_{v,t}$, is computed as $\rho^2_{mv}(\lambda) = \frac{|f_{mv}(\lambda)|^2}{f_{mm}(\lambda) f_{vv}(\lambda)}$, where $f_{mm}(\lambda)$ and $f_{vv}(\lambda)$ are the spectral densities of $w_{m,t}$ and $w_{v,t}$, and $f_{mv}(\lambda)$ the cross-spectral density between $w_{m,t}$ and $w_{v,t}$. These spectral and cross-spectral densities are the components of the spectral density matrix of the bivariate regime factor $w_t = (w_{m,t}, w_{v,t})'$ given by $F_w(\lambda) = A^{-1}(e^{i\lambda}) F_v(\lambda) A^{-1}(e^{i\lambda})^*, \quad \lambda \in [-\pi, \pi]$ where $F_v(\lambda)$ is the $(2 \times 2)$ spectral density matrix of the regime factor innovations $v_t$, and $*$ denotes the adjoint operator. Due to the iid assumption on $v_t = (v_{m,t}, v_{v,t})'$, we have $F_v(\lambda) = P_{vv}$ for all $\lambda \in [-\pi, \pi]$.

(see Figure A.6 of the online appendix). The key takeaway from these figures is that the relation between the mean and volatility factors is time-varying. While the correlation is negative across all rolling windows, the correlation drops to about half during the above-mentioned episodes. In addition, as the coherence illustrates, the correlation is larger for higher frequencies (periods of less than 3 months) thank for the higher frequencies.

## 3.2 Endogenous Feedback, Dynamic Interactions, and Exogenous Regime Switching

To evaluate the in-sample fit of the unsynchronized mean-volatility regime switching model, we first test the null hypothesis of no endogenous feedback from the shock to oil price changes at time $t$ to the regime switching in mean and volatility regime factors at time $t+1$ (i.e., we test the null $\rho_{u,v_m} = \rho_{u,v_v} = 0$). We allow the mean and volatility regime factors to be contemporaneously correlated and interact dynamically through the autoregressive parameters in the matrix $A$ and the correlation parameter $\rho_{v_v,v_m}$. The likelihood ratio test for this hypothesis equals 80.46, thus allowing us to reject the null at a 1% significance level and providing strong evidence that endogenous feedback plays an important role in modeling the dynamics of crude oil prices.

The estimated autoregressive parameters –the off-diagonal elements of the matrix $\boldsymbol{A}$–, reveal a significant dynamic interaction between the mean and volatility regime factors. Estimation results reported in the second panel of Table 1 lead us to reject the null that the previous volatility factor does not influence the current mean factor ($a_{mv} = -0.024$, and statistically significant) as well as the null of no dynamic interaction from the mean factor to the volatility factor ($a_{vm} = -2.177$, and statistically significant). Our estimate of the contemporaneous correlation between the innovations $v_m$ and $v_v$ net of the contribution of the past oil price change innovation $u_{t-1}$ ($\rho_{vv\cdot u} = \rho_{v_v,v_m} - \rho_{v_m,u}\rho_{v_v,u}$) equals $-0.339$. This indicates a negative contemporaneous comovement between the mean and volatility factors, even after netting out the effect of past innovations in crude oil prices. Estimation results not reported herein, but available from the authors upon request indicate that: (a) the model where mean and volatility switching occurs in a synchronized manner as in (3), is rejected in favor of the unsynchronized model (1); (b) our results are robust to estimating the model after excluding the three largest outliers (March, April and May of 2020); and (c) our conclusions are unchanged when we estimate the model using Brent instead of WTI.

### 3.3 Forecasting Evaluation

In this section we compare the forecasting ability of the one- and two-factor models with two alternative models widely used in the empirical regime switching (RS) literature: the standard RS model driven by an exogenous Markov chain (hereafter MKCH) and the RS model with time-varying transition probabilities (hereafter TVTP). As suggested in Diebold et al. (1994), the transition probabilities in the TVTP model evolve as logistic functions of a predetermined variable $z_t$ with parameters that depend upon the regime $s_{t-1}$ at time $t-1$, i.e., $1/\left(1+e^{-\left(\alpha_{s_t}+\beta_{s_t}z_t\right)}\right)$. Regarding the predetermined variables, we consider three alternative TVTP models. One where $z_t$ is given by lagged money supply (M1), another where $z_t$ is the lagged inflation (CPI),[22], and a last one where the lagged dependent variable (WTI) drives the transition probability.

Recall that our one-factor model is given by

$$y_t = \sigma(s_t)u_t \tag{5}$$

In this simple model the volatility process $\sigma(s_t)$ switches over time between high and low levels so that $s_t = 1$ corresponds to the high volatility regime and $s_t = 0$ to the low volatility regime. We will denote this volatility regime switching model by VOL.

The volatility in the next period, $\sigma(s_{t+1})$, is regime-dependent, and therefore we forecast its expectation over two possible regimes as

$$\mathbb{E}\left[\sigma\left(s_{t+1}\right) \mid \mathcal{F}_t, z_t\right] = \sigma(1)p\left(s_{t+1}=1 \mid \mathcal{F}_t, z_t\right) + \sigma(0)p\left(s_{t+1}=0 \mid \mathcal{F}_t, z_t\right) \tag{6}$$

Similarly, for the two-factor model (see equation 1), we forecast the conditional volatility as a weighted average of the high and low volatility with weights given by the probabilities of each of the four states (i.e., low mean-high volatility, low mean-low volatility, high mean-high volatility, high mean-low volatility). That is,

$$\mathbb{E}\left[\sigma\left(s_{v,t+1}\right) \mid \mathcal{F}_t, z_t\right] = \sigma(1)p\left(s_{v,t+1}=1 \mid \mathcal{F}_t, z_t\right) + \sigma(0)p\left(s_{v,t+1}=0 \mid \mathcal{F}_t, z_t\right) \tag{7}$$

where

$$p\left(s_{v,t+1}=1 \mid \mathcal{F}_t, z_t\right) = \sum_{s_{m,t+1}} p\left(s_{v,t+1}=1, s_{m,t+1} \mid \mathcal{F}_t\right)$$
$$p\left(s_{v,t+1}=0 \mid \mathcal{F}_t, z_t\right) = \sum_{s_{m,t+1}} p\left(s_{v,t+1}=0, s_{m,t+1} \mid \mathcal{F}_t\right)$$

[22] The choice of these two predetermined variables is driven by related work (not reported here but available from the authors upon requests), which suggests that the latent factors in our unsynchronized model are most highly correlated with money supply and prices among a large set of macro, financial and oil market variables.

We will denote this unsynchronized regime switching model by UNSYNCH.

To estimate the expected volatility given in (6) and (7), for the VOL and UNSYNCH models respectively, we compute the regime probabilities from the prediction step of the CCP filter. For the TVTP models the transition probabilities, $p(s_{t+1} = 1|F_t, z_t)$ and $p(s_{t+1} = 0|F_t, z_t)$ are computed using the logistic function specified above. Estimates of the high and low volatility levels, $\sigma(1)$ and $\sigma(0)$, are also computed from the respective models. We use five-, ten-, and thirty-year rolling-window samples to construct the forecasts, compute the root-mean-square-error (RMSE) for each model, and compare the out-of-sample performance.

Results reported in Table 2 indicate that our volatility (VOL) model outperforms all alternatives –the conventional Markov switching model with constant transition probabilities (MKCH), and the three regime switching models with time-varying transition probabilities (TVTP-M1, TVTP-CPI, TVTP-WTI)– in terms of RMSE. For ease of interpretation the table also reports the ratio of, the RMSE for the alternative models to the VOL model. As the table reports, for all rolling window schemes the relative performance of the volatility switching model with endogenous feedback, VOL, is superior.

The last column of Table 2 reports the RMSE for the UNSYNCH model for the ten- and thirty-year rolling window schemes. The out-of-sample performance of the UNSYNCH model is worse than that of the VOL model. This result is perhaps not surprising. First, the model with unsynchronized switching in the mean and volatility is more difficult to estimate and will require a longer sample than the simple endogenous volatility switching model to allow for better in-sample fit (Hence, our choice of using larger window sizes). Indeed, the relative performance of the UNSYNCH model improves when we increase the window size from ten to thirty years with the UNSYNCH model going from being the least accurate to being slightly more accurate than all TVTP models but TVTP-WTI. Yet, the larger the estimation window, the greater the probability to incur into misspecification due to nonlinearities, structural changes, etc. Thus, if the object is forecasting out-of-sample, the research might be better served by using the simpler VOL model. All in all, while the more complicated UNSYNCH model has a better in-sample fit, the simpler VOL model produces, on average, more accurate forecasts.

### 3.4 Lessons from Four Episodes of High Volatility and Low Mean

What lessons can oil investors and policy makers derive from a model that allows the regime to evolve in an endogenous manner? At a first glance, gauging the policy implications or empirical relevance of our estimates might seem cumbersome, especially given the multiplicity of states and transition probabilities. We focus on four historical episodes that an analyst or policy maker may have considered as particularly risky for oil markets given the change in the WTI exhibited a low mean and high volatility: the Persian Gulf war, the 2008 financial crisis, the oil price collapse of 2014 and the Covid-19 pandemic. Doing so allows us to distill a clear message: during periods of turmoil, the transition probabilities provide a more realistic assessment of the likelihood of remaining in a known low mean-high-volatility regime than the constant transition probability derived from an exogenous model.

More specifically, we classify a month as exhibiting a low mean if the monthly change in the oil price is negative and lower than ten times the average of the rate of growth in the WTI over the previous 32 months from the start of the particular low-mean-high-volatility regime. Similarly, we classify a month as exhibiting high volatility if the volatility is more than twice the average volatility computed over the 32-month period ending at the start of the given low-mean-high-volatility regime. Months that fall in both categories are classified as low-mean-high-volatility. Note that the source of the disruptions that led to heightened volatility and low mean was distinct and varied across the four episodes, thus providing a good way to illustrate the empirical relevance of our results. Of particular interest is the sample covering the Covid-19 pandemic, which resulted in the oil market experiencing an "all-time high volatility" according to U.S. Energy Information Administration, 2020.

Figure 6 plots the time-varying transition probabilities from a low mean-high volatility to a low-mean-high volatility (hereafter LH-to-LH) regime and the WTI. The top left panel depicts the evolution during the Persian Gulf War. Recall from Figure 1 that the WTI and its volatility increased dramatically when Iraq invaded Kuwait in August 1990; yet it later decreased reaching a trough in March 1991. In turn, the volatility increased with the onset of the war and remained high until Iraq accepted the terms of the cease-fire agreement (March 3, 1991). As illustrated by the black solid line, the LH-to-LH transition probability increased with Iraq's invasion of Kuwait and remained relatively high until the U.S. and allied forces entered Kuwait at the end of February 1991. Contrast the time-varying transition probability with the transition probability obtained from the exogenous switching model denoted by the red dashed line. In the exogenous model, the transition probability remains constant at 0.524. Our time-varying transition probabilities are considerably lower, although they rise throughout the war.

FIGURE 6 HERE

Three additional episodes of low-mean and high-volatility are depicted in Figure 6. The top right panel illustrates the financial crisis of 2008, the bottom left panel corresponds to the oil price collapse of 2014, and the right bottom panel depicts the Great Shutdown resulting from the Covid-19 pandemic. In all cases, the time-varying nature of the transition probability estimated from the endogenous model stands in sharp contrast with the constant transition probability obtained from the exogenous model. The LH-to-LH transition probability increases with the decline in crude oil price growth and stays high until the monthly oil price starts to increase. However, whereas the time-varying probability from our endogenous model remained below the constant probability from the exogenous model for most of the sample, it did not for two episodes. First, during the financial crisis the time-varying transition probability was very close to the probability from the exogenous model. Second, and more notably, in March and April of 2020 when shutdowns and border closures were implemented in the wake of the Covid-19 pandemic, the time-varying transition probability from the endogenous model surpassed the constant transition probability from its exogenous counterpart and then quickly declined in May 2020. The latter coincided with expectations of improved global demand for crude oil and an agreement between OPEC and its allies to cut production. Clearly, transition probabilities derived from an endogenous regime switching model provide a more realistic assessment of the likelihood to remain in the known low mean-high volatility regime. In other words, a lesson to be learned from these four episodes is that not all periods of low mean and high volatility should convey the same degree of "fear" to economic agents.[23]

# 4 Oil Market Uncertainty, Inflation and Inflation Expectations

On the one hand, recent empirical literature on the impact of uncertainty shocks on inflation suggests heightened economic uncertainty acts as a negative demand shock dampening inflation and economic activity. Caggiano et al. (2014), Fernández-Villaverde et al. (2015), Leduc and Liu (2016), Basu and Bundick (2017) and Oh (2020) find a negative effect of uncertainty on inflation. On the contrary, Mumtaz and Theodoridis (2015) find uncertainty has an inflationary effect and Carriero et al. (2018), Katayama and Kim (2018) do not find a significant response of inflation. Work by Oh (2020) indicates the divergence among empirical estimates stems from the use of different sample periods and uncertainty measures

[23] See Figure A.4 of the online appendix for time-varying probabilities and rate of growth in Brent for these four low mean-high volatility periods.

(e.g., financial and macroeconomic uncertainty). On the other hand, fluctuations in oil prices have often been a concern for central bankers as unexpected increases in oil prices have been shown to trigger higher inflation. For instance, in the Press Conference given by the Fed's Chairman Jerome Powell on June 15, 2022 he stated: "[t]he surge in prices of crude oil and other commodities that resulted from Russia's invasion of Ukraine is boosting prices for gasoline and food and is creating additional upward pressure on inflation." Does heightened uncertainty in oil markets result in higher inflation? This section addresses that question.

## 4.1 A Measure of Oil Market Uncertainty

Figure 7 plots the extracted latent volatility factor from our mean-volatility switching model and the Chicago Board Options Exchange (CBOE) Crude Oil ETF Volatility Index (OVX), retrieved from FRED at the Federal Reserve Bank of St. Louis. The OVX measures the 30-day implied volatility of crude oil prices and is computed using fluctuations of the prices of financial options for the WTI. The OVX inception dates from May 10, 2007 and thus covers only part of our sample, but is increasingly cited as a measure of expected volatility in oil markets (see e.g., Energy Information Agency 2020). Note that the extracted volatility factor evolves in a manner similar to the OVX, with both series increasing during the financial crisis, the 2014 oil price collapse and the Great Shutdown. Fluctuations in the OVX somewhat leads fluctuations in the extracted volatility factor as one would expect given that the OVX is a measure of near-term price changes in the WTI. The extracted latent factor dropped faster than the OVX in 2020, suggesting that the increase in oil market volatility was shorter lived than the markets had originally expected. This comparison suggests that the extracted latent volatility factor could be employed as an alternative measure of the overall risk or stress in the oil market.[24]

## 4.2 Comparison with Existing Uncertainty Measures

How does our measure of oil market uncertainty compare to other measures of uncertainty? Table 3 reports the correlations among various measures of uncertainty and the volatility regime factor, $w_v$. The former comprises: (a) the one- and three-month ahead financial and economic uncertainty measures proposed by Jurado, Ludvigson, and Ng (2015), $JLN_{F1}$, $JLN_{E1}$, $JLN_{F3}$, $JLN_{E3}$ respectively; (b) the Chicago Board Options Exchange (CBOE) S&P 100 Volatility Index, VXO; and (c) Baker et al. (2016) U.S. Economic Policy Uncertainty,

[24] Similar results have been found by Chang et al. (2017) for the VIX and the volatility latent regime factor extracted from the U.S. excess market returns.

EPU, index constructed using newspaper coverage of policy related issues. Not surprisingly, as the table shows, the JLN financial uncertainty indices are very highly correlated with the JLN economic uncertainty indices (0.99 in all cases) as they only differ for the later years in the sample. Their correlation with the VXO is high (around 0.8), but smaller with $w_v$ (less than 0.27). The correlation of the EPU with $w_v$ equal 0.22 and ranges between 0.26 and 0.42 for the $JLN$ indices.

TABLE 2 HERE

Figure 8 plots the volatility regime factor, $w_v$ , against four measures of uncertainty. In panels (a) and (b), respectively, we compare $w_v$ to the one- and three-month ahead financial uncertainty measures proposed by Jurado, Ludvigson, and Ng (2015).[25] As the figure illustrates, periods of heightened financial uncertainty such as the Great Recession and the Covid-19 pandemic coincided with periods of higher uncertainty in oil markets. Yet, uncertainty in oil markets was markedly higher during the Persian Gulf war and the mid-2010s. A similar pattern is observed in panel (c) when comparing the latent factor with the CBOE S&P 100 Volatility Index, VXO. Finally, panel (d) contrast the volatility regime factor and the Economic Policy Uncertainty (EPU) index computed by Baker et al. (2015). Again, in this case, while some of the peaks coincide, the volatility regime factor reflects 'risk' factors not captured by the EPU.

FIGURE 8 HERE

To summarize, as Table 3 and Figure 8 illustrate, while the volatility regime factor commoves with other uncertainty indices, it captures a different type of uncertainty: oil market risk.

## 4.3 The Effect of Oil Market Uncertainty on Inflation, Inflation Expectations, and Disagreement

To assess the impact of oil market uncertainty, as measured by our mean and volatility regime factors, on inflation, inflation expectations and disagreement we estimate a five-variable vector autoregressive (VAR) model given by

$$B_0 y_t = c + B_1 y_{t-1} + B_2 y_{t-2} + ... + B_p y_{t-p} + \epsilon_t \tag{8}$$

[25]Given that the correlation between the JLN financial and economic uncertainty indices is 0.99 we restrict the plots to the financial indices.

where $c$ is an intercept, $y_t$ is a $5 \times 1$ vector that includes the volatility regime factor ($w_v$), the mean regime factor ($w_m$), the CPI inflation rate ($\pi_t$), the median expected inflation ($\pi_t^{i,exp}$), and the inflation expectations inter-quartile range as a measure of disagreement($\pi_t^{i,dis}$). We use a conservative lag order $p = 4$ (see Kilian and Lütkepohl 2017) leads to similar results. $B_i, i = 0, ..., p$ denote $5 \times 5$ coefficient matrices and $\epsilon_t$ represent a vector of mutually uncorrelated i.i.d. structural shocks. The model is partially identified in that only the first two structural shocks are identified by zero-restrictions on the $B_0$ matrix. The impulse response estimates of interest are invariant to the identification of the remaining structural shocks (see Christiano, Eichenbaum and Evans 1999) and the responses of the last three variables are robust to changing the order of the latent factors.

We first estimate a VAR for the short-run (one-year), $i = 1yr$, inflation expectations obtained from the Survey of Professional Forecasters (SPF) and show that the results are robust if we measure expectations using the University of Michigan (UofM) Consumer Survey. We then re-estimate the VAR by replacing the short-run inflation expectations and disagreement with their medium (five-year), $i = 5yr$, and long-run (ten-year), $i = 10yr$, counterparts from the SPF.[26] The volatility and mean regime factors are those extracted from our benchmark model. Given that data from the SPF are quarterly, we aggregate the monthly factors and the UofM data by taking the average over the quarter.

Figure 9 depicts the effect of shocks to the volatility and mean regime factors when we use the SPF to measure short-run inflation expectations. The solid line represents the impulse response and the light and dark shaded areas represent 68 and 90 percent confidence bands, respectively, computed by bootstrap. An increase in oil market volatility factor has a negative effect on inflation on impact that vanishes quickly, a negative but short-lived effect on short-run inflation expectations, and a positive non persistent effect on disagreement. As for the mean factor, an increase in the mean of oil price changes leads to an increase in inflation, only a marginal but persistent increase in short-run inflation expectations and an increase in disagreement.

FIGURE 9 HERE

Figure 10 reports impulse response functions obtained when we use the UofM consumer survey. Qualitatively, the results for the response of inflation, inflation expectations and disagreement are similar to those obtained using the SPF. However, two differences stand

[26] Estimation results are robust to using the 5-10 year inflation expectations from the University of Michigan survey.

out. First, the impact of increased oil market uncertainty results in disagreement among households about inflation expectations that is an order of magnitude higher than estimated for the SPF. Second, an increase in the mean regime factor leads to a significant increase in household's short-run inflation expectations that is somewhat more persistent than estimated for the SPF. While estimation results using monthly data for the UofM survey suggests the deflationary effect of oil uncertainty shocks might be insignificant, the finding of a positive (and persistent) impact on households' inflation expectations and disagreement is robust to using higher frequency data.

FIGURE 10 HERE

How do medium and long-run inflation expectations respond to changes in the mean and volatility factors? Could shocks to oil price changes pose a challenge to central banks by causing the inflation anchor to drift? To answer these questions we now focus on the responses of medium (5-year) and long-run (10-year) inflation expectations and disagreement among professional forecasters (SPF) to innovations in the volatility and mean regime factors. For the sake of brevity, we restrict our discussion to these expectation responses, but do plot the inflation responses in the figures.

The top panels of Figures 11 and 12 show that an increase in oil market volatility regime factor leads to a significant decline in medium inflation expectations and no significant impact in long-run inflation expectations. As mentioned earlier, this response is consistent with increases in oil market volatility working as a negative demand shock. Furthermore, higher volatility factor results in increased disagreement in the medium-run and a smaller, albeit only marginally significant, increase in the long-run.

FIGURE 11 HERE

Regarding shocks to the mean factor, the bottom panels of Figures 11 and 12 illustrate how increases in the mean of oil price changes result in raises in the 5- and 10-year inflation expectations. We note that whereas the effect of shocks to the volatility latent factor on inflation expectations and disagreement appear to dissipate in the long-run, the effect of shocks to the mean factor do not. Note that the magnitude of the increase in inflation expectations in the short, medium and long-run is similar (about a tenth of the unexpected increase in the mean factor) and it retains statistical significance at 10 years. Similarly, on impact, disagreement at 5 and 10-years increases and remains significant for a few quarters after the shock.

FIGURE 12 HERE

All in all, estimation results presented in this section illustrate how the extracted latent factors may be used by policy makers and analysts as a way to measure uncertainty in oil markets.[27] Moreover, we show that oil market uncertainty results in increased disagreement about future inflation among professional forecasters and, especially, among households. While the effect on disagreement is statistically significant at all horizon, its magnitude falls as the horizon expands.

Why does the increase in disagreement matter? As Reis (2021) posits, disagreement about future inflation could be an early signal of a shift in the inflation anchor and "[if] expectations persistently change, then the anchor is adrift; if they differ from the central bank's target, the anchor is lost". Now, could oil price fluctuations contribute to movements in the inflation anchor? Our estimation results suggest that because increases in the volatility regime factor are associated with a decrease in the mean regime factor of oil price changes, they do not appear to pose a risk for inflation anchoring in the medium and long run. Similarly, increases in the mean factor of oil price changes could result in higher short-run inflation expectations and disagreement, thus only posing a risk to inflation anchoring in the short-run.

# 5 Conclusions

We employed an endogenous regime switching model to study the behavior of the crude oil prices. To gain some intuition regarding the importance of allowing for endogenous switching, we started our investigation using a volatility switching model. We built on the model by allowing for unsynchronized switching in the mean and volatility of crude oil prices.

Forecasting comparison exercises show that our regime switching models with endogenous feedback compare well with other regime switching models with time varying transition probabilities driven by exogenous or past endogenous variables. We note that our endogenous regime switching models do not require specifying the specific form of endogeneity, which is a convenient advantage in practical implementation. Using this model we constructed two new measures of oil market uncertainty (the volatility and mean regime factors). We

[27] While we opt for a smaller scale VAR, we note that the results are robust to rotating in variables commonly used in the study of inflation: the unemployment rate, the effective federal funds rate, and two measures of economic activity (i.e., the rate of growth in the industrial production index or the Chicago Fed National Activity Index) in the SVAR.

subsequently estimate a SVAR to investigate the effect of these novel uncertainty shocks on actual inflation, expected inflation, and disagreement.

Four key results are derived from our paper. First, we demonstrated that a model that allows the processes that govern the switching between volatility and mean regimes to evolve in an endogenous manner produces a better in-sample fit than an exogenous regime switching model. Moreover, conditional on knowing the regime, an endogenous regime switching model provides useful information regarding the time-varying nature of the volatility and, hence, could be useful in assessing risk.

Second, we showed that the volatility regime switching model with endogenous feedback does a good job in out-of-sample forecasts. Indeed, based on the RMSE, it can outperform two alternative regime switching models commonly used in the literature.

Third, we showed that the extracted latent volatility factor from the unsynchronized mean-volatility switching model can be used to gauge risk in crude oil markets. In fact, while it commoves with measures of economic policy uncertainty, it appears to capture a different type of uncertainty than that embedded in existing macroeconomic and financial uncertainty measures.

Finally, using a VAR model we illustrated how the extracted latent regime factors can provide useful information to policy makers as unexpected increases in the volatility factor lead to declines in CPI inflation and inflation expectations, but also results in increased disagreement among professional forecasters and, especially, households regarding short-run expectations of future inflation. We note that increases in the volatility and the mean factors could be used by central bankers as an early warning sign that the inflation anchor might drift, especially in the short run.

To summarize, estimation results presented in this paper suggest the use of endogenous regime switching models could be useful for policy analysts and economic agents interested in understanding fluctuations in crude oil prices. In particular, the extracted volatility regime factor may be used as a measure of risk or stress in oil markets. Tracking the evolution of this latent factor can not only serve as an early warning signal of oil market risk, but also aid in understanding the evolution of inflation, inflation expectations and disagreement during periods of high oil market uncertainty.

# 6 References


Baker, S. R., Bloom, N. and Davis, S. J. (2016). "Measuring economic policy uncertainty," *The Quarterly Journal of Economics*, 131(4), 1593-1636.

Basu, S., and Bundick, B. (2017). "Uncertainty shocks in a model of effective demand," *Econometrica*, 85(3), 937-958.

Baumeister, C. and Hamilton, J. D. (2019). "Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks," *American Economic Review*, 109(5), 1873-1910.

Baumeister, C., Korobilis, D. and Lee, T. K. (2020). "Energy markets and global economic conditions," *The Review of Economics and Statistics*, 9, 1-45.

Benigno, G., Foerster, A., Otrok, C. and Rebucci, A. (2020). "Estimating Macroeconomic Models of Financial Crises: An Endogenous Regime-Switching Approach," NBER Working Paper NBERw26935 available at https://www.nber.org/papers/w26935.

Binder, C. C., 2018. "Inflation expectations and the price at the pump," *Journal of Macroeconomics*, 58, 1-18.

Bjørnland, H. C., Larsen, V. H. and Maih, J. (2018). "Oil and macroeconomic (in) stability," *American Economic Journal: Macroeconomics*, 10(4), 128-51.

Caggiano, G., Castelnuovo, E., and Groshenny, N. (2014). "Uncertainty shocks and unemployment dynamics in US recessions," *Journal of Monetary Economics*, 67, 78-92.

Carriero, A., Clark, T. E. and Marcellino, M., 2018. "Measuring uncertainty and its impact on the economy," *Review of Economics and Statistics*, 100(5), 799-815.

Chang Y., Choi Y. and Y. Park, (2017). "A New Approach to Model Regime Switching," *Journal of Econometrics*, 196, 127-143.

Chang, Y., Kwak, B., Park, J. Y. and Qiu, S. (2021). Regime Switching Models with Multiple Dynamic Factors. Mimeograph, Indiana University.

Chang, Y., Kwak, B. and Qiu, S. (2021). "U.S. Monetary and Fiscal Policy Regime Changes and Their Interactions," Mimeograph, available at https://economics.indiana.edu/documents/pap monetary-and-fiscal-policy-regime-changes-and-their-interactions.pdf

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), pp.197-232.

Christiano, L. J., Eichenbaum, M. and Evans, C. L. (1999). "Monetary policy shocks: What have we learned and to what end?," Handbook of Macroeconomics, in: J. B. Taylor & M. Woodford (ed.), *Handbook of Macroeconomics*, edition 1, volume 1, chapter 2, pp.65-148 Elsevier.

Davig, T. and Leeper, E. M. (2006). "Endogenous Monetary Policy Regime Change," in: Reichlin, L. and West, K.D. (Ed.), *NBER International Seminar on Macroeconomics 2006*. University of Chicago Press, pp. 345–391.

Di Sanzo, S. (2018). "A Markov switching long memory model of crude oil price return volatility," *Energy Economics*, 74, 351-359.

Diebold. F., Lee, J. H. and Weinbach, H. (1994). Regime switching with time-varying transi-tion probabilities. In C. Hargreaves (Ed.), *Nonstationary Time Series Analysis and Cointe-gration, Advanced Texts in Econometrics*, pp. 283–302. Oxford University Press.

Fattouh, B., Kilian, L. and Mahadeva, L. (2013). "The role of speculation in oil markets: What have we learned so far?," *The Energy Journal*, 34(3), 7-33.

Fernández-Villaverde, J., Guerrón-Quintana, P., Kuester, K. and Rubio-Ramírez, J. (2015). "Fiscal volatility shocks and economic activity," *American Economic Review*, 105(11), 3352-84.

Filardo A. J. (1994). "Business-Cycle Phases and their Transitional Dynamics," *Journal of Business & Economic Statistics*, 12(3), 299–308.

Hamilton, James D. (1989). "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle," *Econometrica* 57, no.2, 357–84.

Hamilton, J. D. (2010). "Regime switching models," in *Macroeconometrics and Time Series Analysis* (pp. 202-209), Palgrave Macmillan, London.

Herrera, A. M., and Rangaraju, S. K. (2020). "The effect of oil supply shocks on US economic activity: What have we learned?," *Journal of Applied Econometrics*, 35(2), 141-159.

Herrera, A. M., Hu, L. and Pastor, D. (2018). "Forecasting crude oil price volatility," *International Journal of Forecasting*, 34(4), 622-635.

Holm-Hadulla, F. and Hubrich, K. (2017). "Macroeconomic implications of oil price fluctuations: A regime-switching framework for the euro area", Working Paper, European Central Bank.

Hubrich, K. and Waggoner, D. (2021). "The Transmission of Financial Shocks and

Leverage of Banks: An Endogenous Regime Switching Framework," Working Paper available at https://papers.ssrn.com/sol3/papers.cfm?abstract$_i d = 3754617$.

Jurado, K., Ludvigson, S. C. and Ng, S. (2015). "Measuring uncertainty," *American Economic Review*, 105(3), 1177-1216.

Katayama, M. and Kim, K. H. (2018). "Uncertainty shocks and the relative price of investment goods," *Review of Economic Dynamics*, 30, 163-178.

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), 1053-69.

Kilian, L., and Lütkepohl, H. (2017). *Structural Vector Autoregressive Analysis* (Themes in Modern Econometrics). Cambridge: Cambridge University Press.

Kilian, L. and Zhou, X. (2022). "The impact of rising oil prices on US inflation and inflation expectations in 2020–23," *Energy Economics*, 113, p.106228.

Leduc, S. and Liu, Z. (2016). "Uncertainty shocks are aggregate demand shocks," *Journal of Monetary Economics*, 82, 20-35.

Mumtaz, H. and Theodoridis, K. (2015). "The international transmission of volatility shocks: An empirical analysis," *Journal of the European Economic Association*, 13(3), 512-533.

Oh, J. (2020). "The propagation of uncertainty shocks: Rotemberg versus Calvo," *International Economic Review*, 61(3), 1097-1113.

Politis, D. N. and Romano, J. P. (1994). "The stationary bootstrap," *Journal of the American Statistical Association*, 89(428), 1303-1313.

Powell, J.(2022). Transcript of Chair Powell's Press Conference, June 15 2022.

Reis, R. (2022). "Losing the inflation anchor," *Brookings Papers on Economic Activity*, 2021(2), 307-379.

Rotemberg, J.J. (1982). "Sticky prices in the United States," *Journal of political economy*, 90(6), 1187-1211.

U.S. Energy Information Administration (2020). "Today in Energy" available at https://www.eia.gov/todayinenergy/detail.php?id=43275

Table 1: **Estimation Results- Regime Switching Models**

| Parameters | Volatility switching | 68% CI | Unsynchronized switching | 68% CI |
|---|---|---|---|---|
| $\tau$ | 5.691 | [2.893, 5.900] | | |
| $\tau_1$ | | | -1.033 | [-1.165, -0.983] |
| $\tau_2$ | | | 9.873 | [8.339, 10.945] |
| $\rho_{u,v}$ | -0.974 | [-0.992, -0.333] | | |
| $\rho_{u,v_m}$ | | | 0.529 | [0.335, 0.540] |
| $\rho_{u,v_v}$ | | | -0.917 | [-0.956, -0.340] |
| $\rho_{v_m,v_v}$ | | | -0.824 | [-0.825, -0.414] |
| $\rho_{v_m v_v \cdot u}$ | | | -0.339 | [-0.609, -0.191] |
| $\alpha$ | 0.971 | [0.912, 0.978] | | |
| $a_{mm}$ | | | 0.133 | [0.060, 0.242] |
| $a_{vm}$ | | | -2.177 | [-2.278, -1.868] |
| $a_{mv}$ | | | -0.024 | [-0.024, -0.024] |
| $a_{vv}$ | | | 0.847 | [0.843, 0.848] |
| $\sigma_0$ | 0.069 | [0.064, 0.072] | 0.059 | [0.056, 0.059] |
| $\sigma_1$ | 0.222 | [0.160, 0.277] | 0.219 | [0.151, 0.257] |
| $\mu_0$ | | | -0.085 | [-0.093, -0.085] |
| $\mu_1$ | | | 0.018 | [0.015, 0.022] |
| $\gamma_1$ | | | 0.084 | [0.084, 0.185] |
| log-likelihood | 455.624 | [433.299, 500.684] | 469.240 | [428.357, 493.959] |

Table 2: **One Step-Ahead Forecast Comparison**

| | VOL | MKCH | TVTP-M1 | TVTP-CPI | TVTP-WTI | UNSYNCH |
|---|---|---|---|---|---|---|
| *Five-year window* | | | | | | |
| RSME | 0.0376 | 0.0389 | 0.0403 | 0.0398 | 0.0394 | |
| Relative RMSE | 100.00 | 103.41 | 106.94 | 105.73 | 104.62 | |
| *Ten-year window* | | | | | | |
| RSME | 0.0418 | 0.0422 | 0.0430 | 0.0433 | 0.0430 | 0.0441 |
| Relative RMSE | 100.00 | 101.59 | 102.88 | 103.67 | 102.83 | 105.50 |
| *Thirty-year window* | | | | | | |
| RSME | 0.0895 | 0.0911 | 0.0928 | 0.0927 | 0.0925 | 0.0927 |
| Relative RMSE | 100.00 | 101.82 | 103.77 | 103.63 | 103.40 | 103.58 |

Notes: VOL indicates our one-factor endogenous volatility switching model, MKCH indicate the conventional Markov switching model with constant transition probabilities, TVTP-M1 and TVTP-CPI indicate the regime switching models with time-varying transition probabilities based on M1 and CPI, respectively, and UNSYNCH denotes our benchmark model with endogenous switching in the mean and variance.

Table 3: **Correlation Among Uncertainty Measures**

| | $w_v$ | $JLN_{F1}$ | $JLN_{F3}$ | $JLN_{E1}$ | $JLN_{E3}$ | VXO | EPU |
|---|---|---|---|---|---|---|---|
| $w_v$ | 1.000 | | | | | | |
| $JLN_{F1}$ | 0.263 | 1.000 | | | | | |
| $JLN_{F3}$ | 0.264 | 0.999 | 1.000 | | | | |
| $JLN_{E1}$ | 0.250 | 0.997 | 0.997 | 1.000 | | | |
| $JLN_{E3}$ | 0.251 | 0.999 | 0.998 | 0.995 | 1.000 | | |
| VXO | 0.215 | 0.810 | 0.805 | 0.812 | 0.807 | 1.000 | |
| EPU | 0.225 | 0.264 | 0.419 | 0.423 | 0.382 | 0.481 | 1.000 |

Notes: $w_v$ denotes the volatility regime factor extracted from our benchmark model; $JLN_{F1}$ and $JLN_{F3}$ denote the one- and three-months financial uncertainty measures from of Jurado, Ludvigson and Ng (2015) whereas $JLN_{E1}$ and $JLN_{E3}$ denote the economic uncertainty measures; VXO is the CBEO volatility index and EPU is the economic policy uncertainty index from Baker et al. (2015).

Figure 1: Evolution of WTI Price

(a) Monthly WTI price and its rate of growth

(b) Monthly WTI volatility

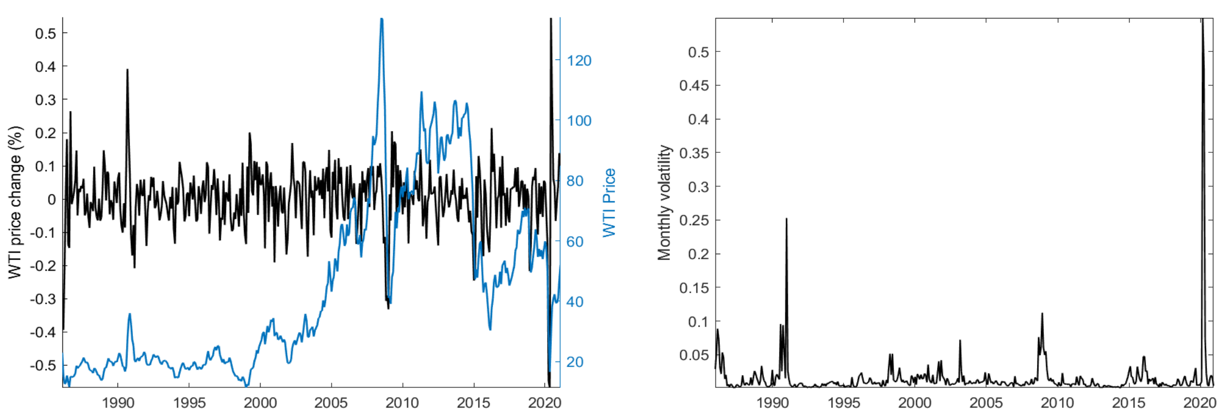


Notes: Panel (a) plots monthly West Texas Intermediate (WTI) prices (blue line) and its rate of growth (black line). Panel (b) plots the monthly volatility of the percentage change in the WTI price.

Figure 2: Extracted Latent Factor: Volatility Switching Model

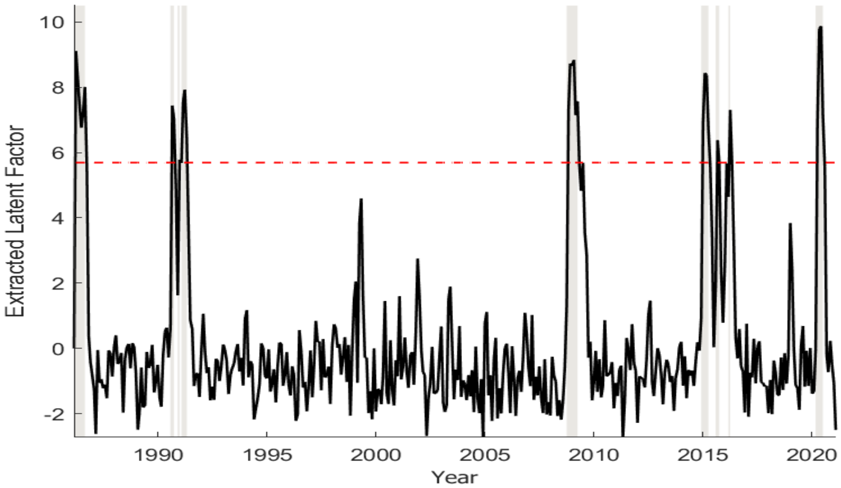


Notes: This figure depicts the extracted latent factor from the volatility switching model (solid black line) and the estimated threshold $\tau$ (dashed red line).

Figure 3: Transition Probabilities: Volatility Switching Model

(a) High to High

(b) Low to High

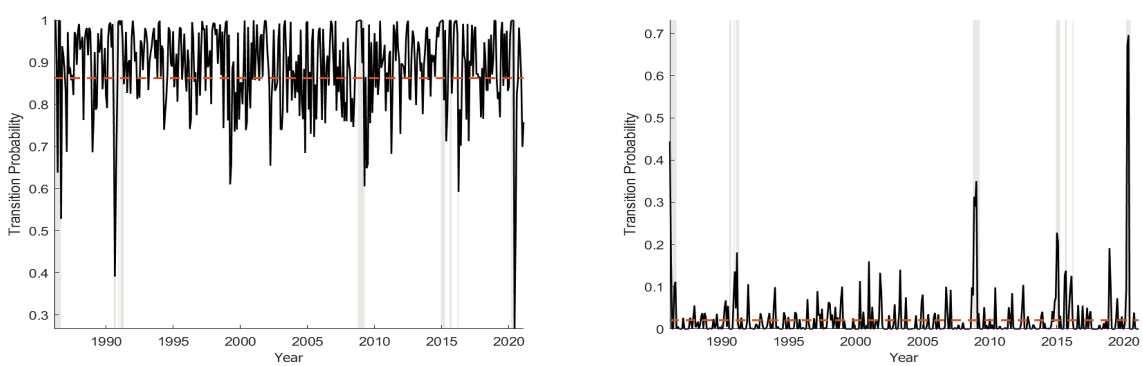


Notes: This figure presents the transition probabilities from the volatility switching model. Panel (a) depicts the transition probability from high to high volatility state: the black solid line corresponds to the time-varying transition probabilities from the endogenous volatility switching model, while the red dashed line represents the constant transition probability from the exogenous model. Similarly, panel (b) depicts the transition probabilities of switching from the low to the high volatility regime.

Figure 4: Extracted Factors - Unsynchronized Mean and Volatility Switching

(a) Mean Factor (b) Volatility Factor

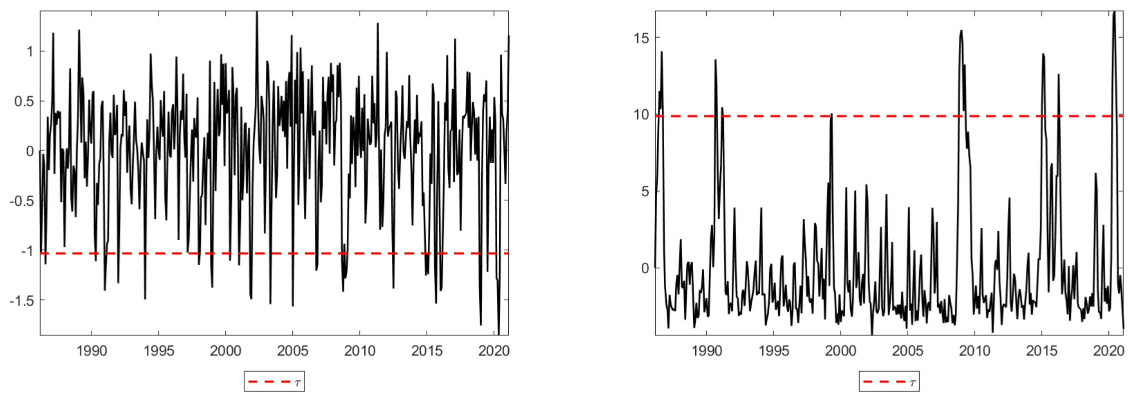


Notes: This figure depicts the extracted latent factors from the unsynchronized mean-volatility switching model. The solid black line in panel (a) corresponds to the mean regime factor and the red dashed line represents the estimated threshold $\tau_m$. Similarly, the solid black line in panel (b) represents to the volatility regime factor and the red dashed line depicts the corresponding threshold $\tau_v$.

Figure 5: Correlation and Coherence between Extracted Latent Factors

(a) Correlation (b) Coherence

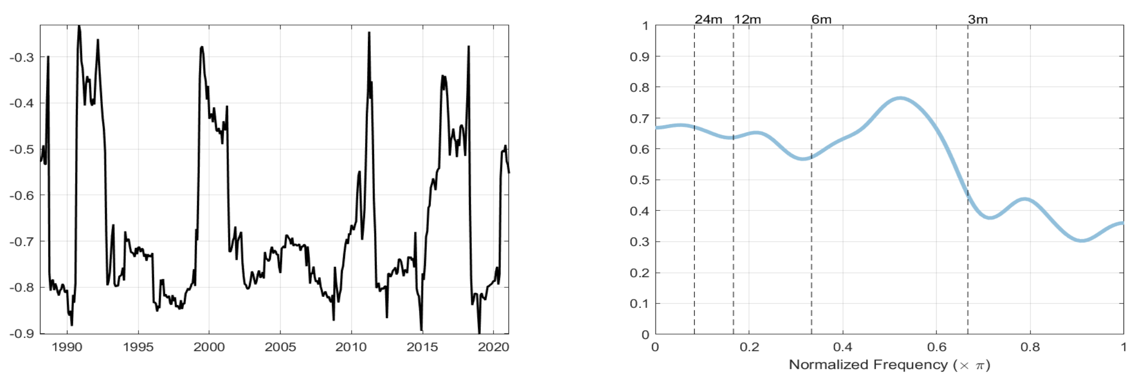


Notes: Panel (a) depicts the correlation between the mean and volatility regime factors computed using a 24-month rolling window. Panel (b) depicts the coherence between the two latent factors computed using the full sample.

Figure 6: LH-to-LH Transition Probabilities and WTI Price Changes

(a) Persian Gulf War

(b) Great Recession

(c) Oil Price Collapse

(d) Covid-19 Pandemic

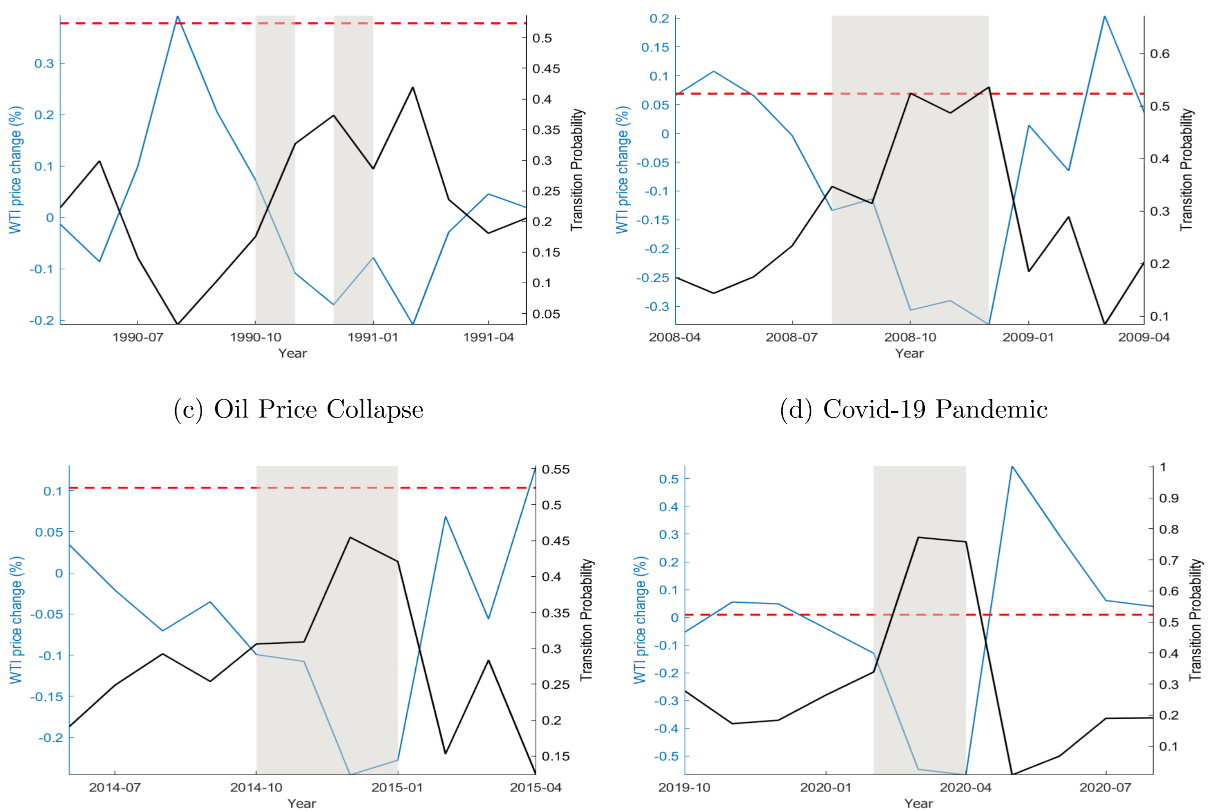


Notes: This figure illustrates the probability to remain in low mean-high volatility state for four episodes of turbulence in oil markets. The black solid line represents the transition probability $\mathbb{P}(s_{m,t} = 0, s_{v,t} = 1|s_{m,t-1} = 0, s_{v,t-1} = 1, y_{t-1})$ estimated from the endogenous unsynchronized switching model; the dashed red line corresponds to the constant transition probability $\mathbb{P}(s_{m,t} = 0, s_{v,t} = 1|s_{m,t-1} = 0, s_{v,t-1} = 1)$ estimated from the exogenous switching model; the solid blue line corresponds to the percentage change in the WTI price. The shaded areas denote periods of low mean and high volatility.

Figure 7: Extracted Volatility Regime Factor from Mean-Volatility Switching Model and CBOE Crude Oil Volatility Index (OVX)

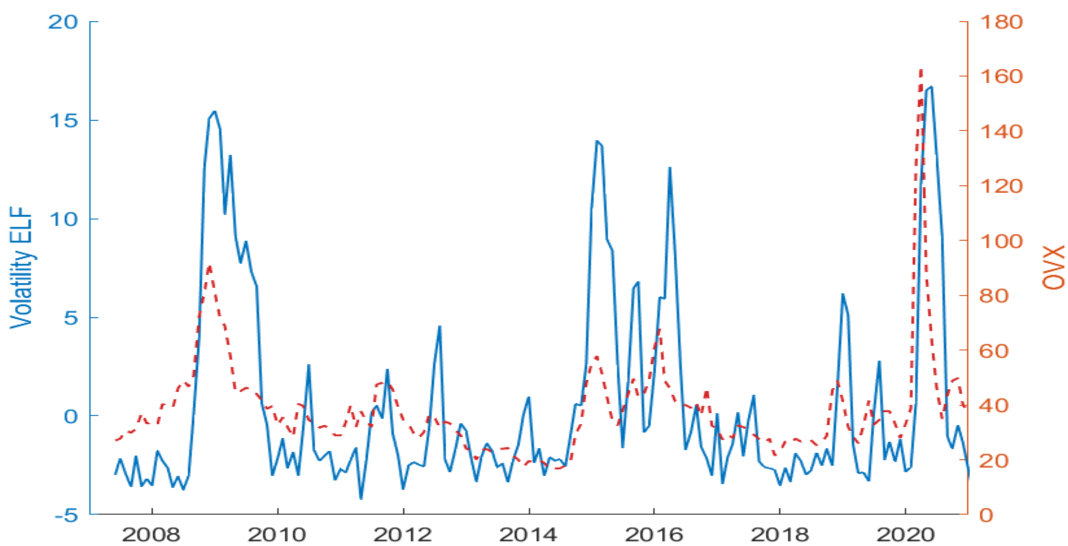


Notes: This figure plots the volatility regime factor extracted from the unsynchronized mean-volatility switching model (solid blue line) and the Chicago Board Options Exchange (CBOE) Crude Oile Volatility Index OVX (dashed red line).

Figure 8: Uncertainty Measures and Regime Factors


Notes: Each panel show the volatility factor estimated from the asynchronized two factor model (blue) and various measures of uncertainty (red).

Figure 9: Response to Latent Factor Shocks - SPF Short-Run Inflation Expectations

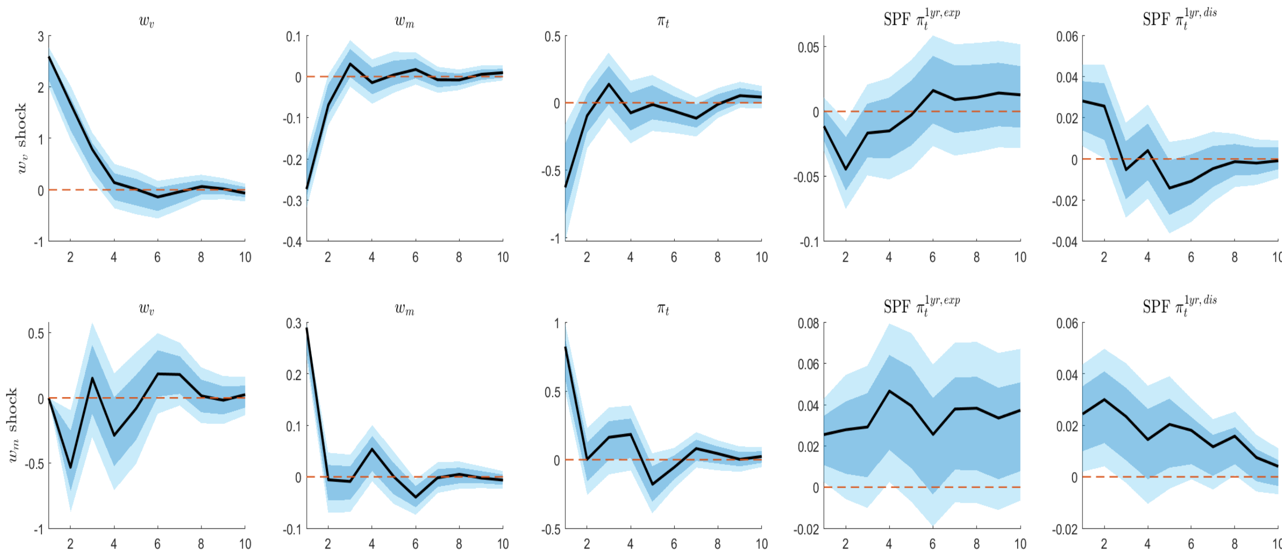


Notes: The top panels illustrate the response to a volatility factor shock. The bottom panels depict the response to a mean factor shock. 68% and 90% confidence bands, denoted by dark and blue shaded areas respectively, are computed by bootstrap.

Figure 10: Response to Latent Factor Shocks - UofM Short-Run Inflation Expectations

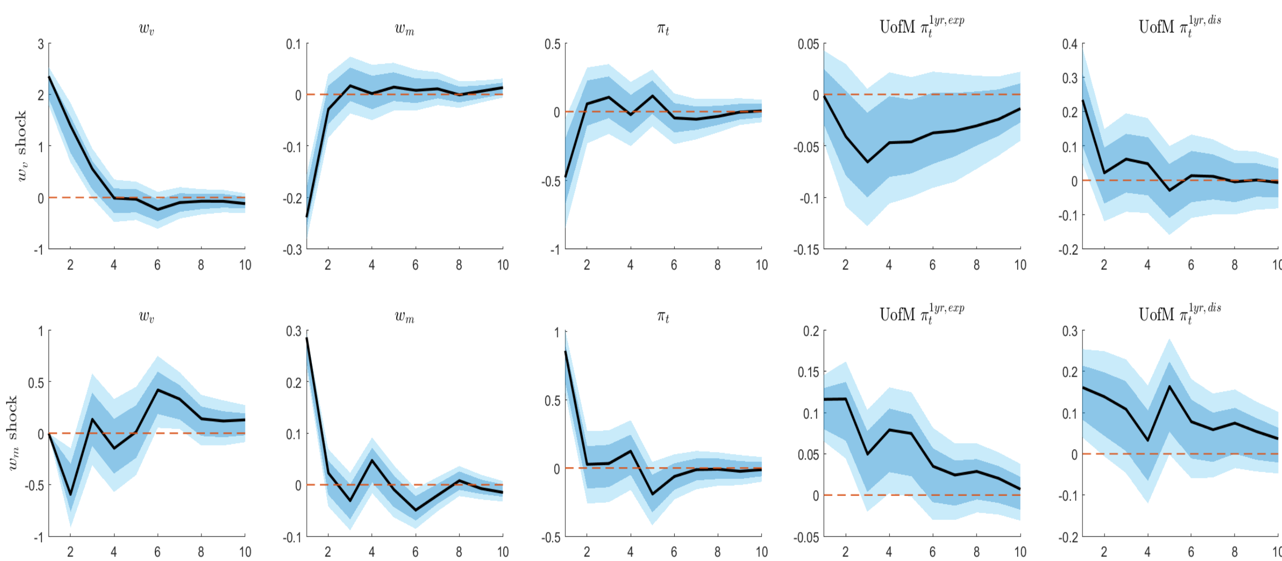


Notes: The top panels illustrate the response to a volatility factor shock. The bottom panels depict the response to a mean factor shock. 68% and 90% confidence bands, denoted by dark and blue shaded areas respectively, are computed by bootstrap.

Figure 11: Response to Latent Factor Shocks - SPF Medium-Run Inflation Expectations

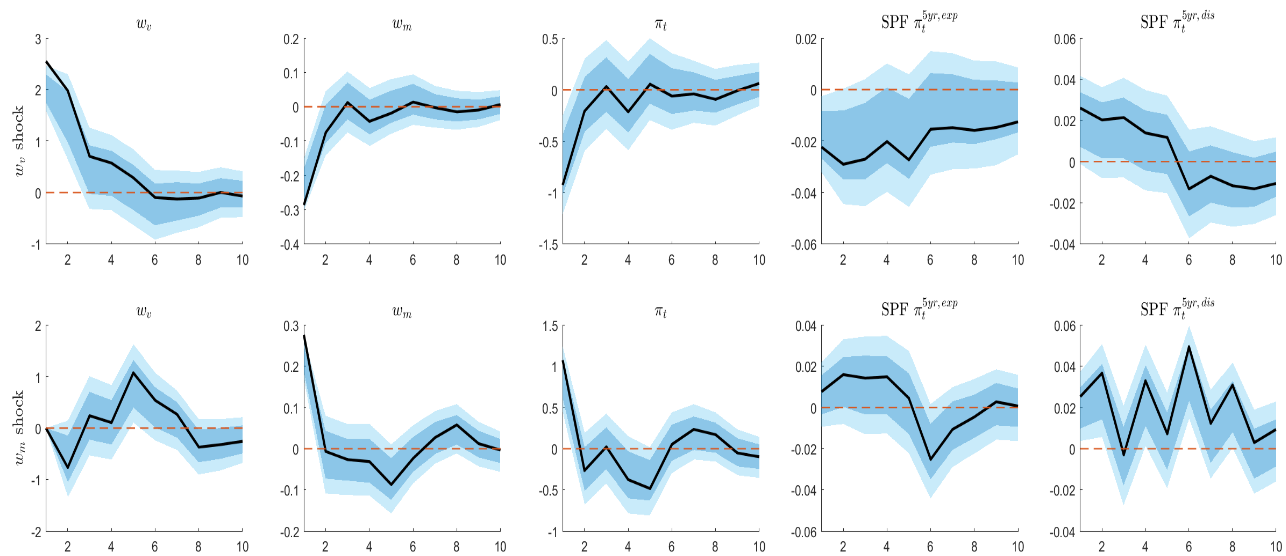


Notes: The top panels illustrate the response to a volatility factor shock. The bottom panels depict the response to a mean factor shock. 68% and 90% confidence bands, denoted by dark and blue shaded areas respectively, are computed by bootstrap.

Figure 12: Response to Latent Factor Shocks - SPF Long-Run Inflation Expectations

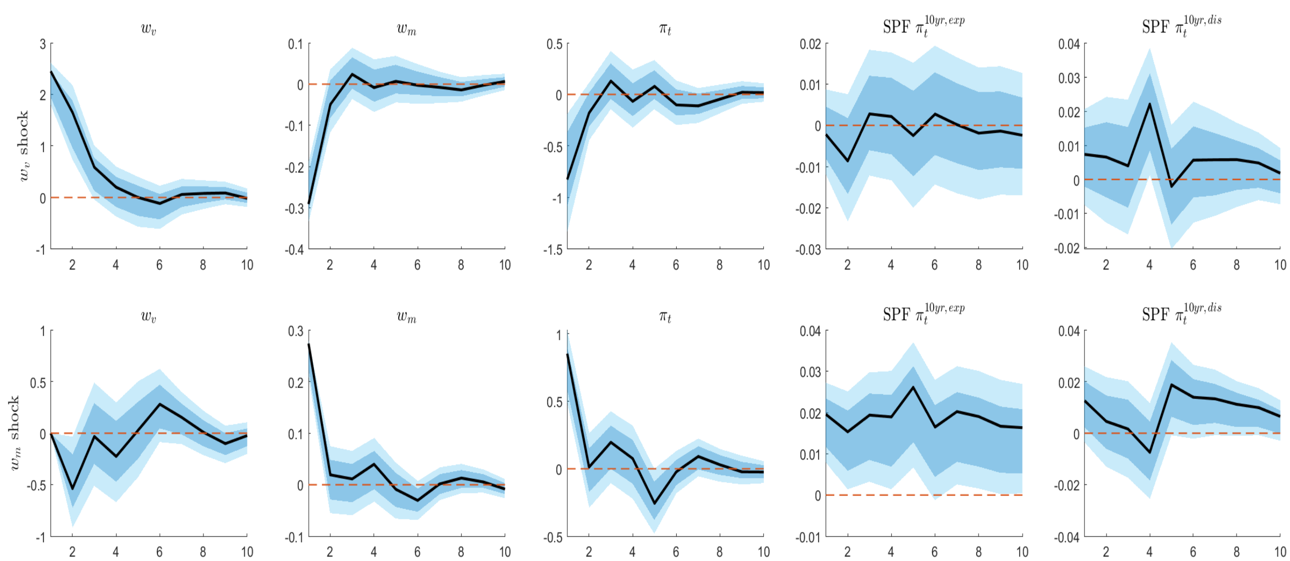


Notes: The top panels illustrate the response to a volatility factor shock. The bottom panels depict the response to a mean factor shock. 68% and 90% confidence bands, denoted by dark and blue shaded areas respectively, are computed by bootstrap.