Research Division
Federal Reserve Bank of St. Louis
Working Paper Series

Do Oil Shocks Drive Business Cycles? Some U.S. and International Evidence

Kristie M. Engemann
Kevin L. Kliesen
and
Michael T. Owyang

Working Paper 2010-007D
http://research.stlouisfed.org/wp/2010/2010-007.pdf

March 2010
Revised March 2011

FEDERAL RESERVE BANK OF ST. LOUIS
Research Division
P.O. Box 442
St. Louis, MO 63166

The views expressed are those of the individual authors and do not necessarily reflect official positions of the Federal Reserve Bank of St. Louis, the Federal Reserve System, or the Board of Governors.

Federal Reserve Bank of St. Louis Working Papers are preliminary materials circulated to stimulate discussion and critical comment. References in publications to Federal Reserve Bank of St. Louis Working Papers (other than an acknowledgment that the writer has had access to unpublished material) should be cleared with the author or authors.

Do Oil Shocks Drive Business Cycles? Some U.S. and International Evidence*

Kristie M. Engemann† Kevin L. Kliesen‡ Michael T. Owyang§

Federal Reserve Bank of St. Louis
P.O. Box 442
St. Louis, MO 63166

First Draft: February 20, 2010
This Draft: January 29, 2011

Abstract

Oil prices rose sharply prior to the onset of the 2007-2009 recession. Hamilton (2005) noted that nine of the last ten recessions in the United States were preceded by a substantial increase in the price of oil. In this paper, we consider whether oil price shocks significantly increase the probability of recessions in a number of countries. Because business cycle turning points generally are not available for other countries, we estimate the turning points together with oil's effect in a Markov-switching model with time-varying transition probabilities. We find that, for most countries, oil shocks do affect the likelihood of entering a recession. In particular, for a constant, zero term spread, an average-sized shock to WTI oil prices increases the probability of recession in the U.S. by nearly 50 percentage points after one year and nearly 90 percentage points after two years. [JEL: C32, E32]

*The authors benefited from conversations with Lutz Kilian, Gert Peersman, and Jeremy Piger. Michelle Armesto, Charles Gascon, and Doug Smith provided research assistance. Views expressed here are the authors' alone and do not reflect the opinions of the Federal Reserve Bank of St. Louis or the Federal Reserve System. All remaining errors are the sole responsibility of the authors.

†kristie.m.engemann@stls.frb.org

‡kliesen@stls.frb.org

§owyang@stls.frb.org

Running Head: Oil Shock and Business Cycle

Corresponding Author:

Kevin Kliesen

Federal Reserve Bank of St. Louis

P.O. Box 442

St. Louis, MO 63166

Fax: (314) 444-8731

Tel: (314) 444-8583

Email: kliesen@stls.frb.org

Keywords: time-varying transition probabilities, probit, recession, oil prices, Gibbs sampler

1 Introduction

Oil price shocks are often thought to be one of the principal drivers of economic fluctuations. Hamilton (2005) argued that nine out of the ten recessions that occurred between 1948 and 2001 were preceded by a rise in oil prices. Most studies, however, assume a proportional relationship between the magnitude of oil price shocks and their effect on economic activity, a notion contrary to the traditional characterization of recessions as being distinct business cycle phases.[1] In more recent business cycle models [e.g., Hamilton (1989)], a hidden Markov variable can be used to reflect expansion or recession. In these models, the transition between phases is governed by a constant probability process that is not influenced by other macroeconomic forces. For the purposes of policy, however, it may be beneficial to determine which forces – if any – influence these transitions.

We consider the possibility that oil price shocks influence the probability that the economy enters recession by estimating a hidden Markov model with time-varying transition probabilities. In our model, increases in (net) oil prices raise the likelihood that the economy transitions from expansion to recession. We estimate the model for a panel of industrialized countries as the effectiveness of oil prices at driving business cycle transitions may not be homogeneous across countries [for example, see Peersman and Van Robays (2009)]. For instance, oil's influence may be a rising function of, say, per-capita energy consumption.

This paper unites two strands of the literature. The first investigates the timing and predictability of business cycle turning points. While the early literature on business cycle turning points was primarily narrative, more recent work has used more formal econometric techniques. Taking the National Bureau of Economic Research (NBER) Business Cycle Dating Committee's turning points as given, Estrella and Mishkin (1998) used a probit model to show that financial market variables can be used to forecast recessions. A large body

[1] More recent studies beginning with Hamilton (1996, 2003, 2005) define a net oil price shock. While the net oil shock is a nonlinear function of prices (a rise in oil prices has an effect on economic activity only if the price exceeds a threshold), its effect is still linear. A small net oil shock has a small effect and a large net oil shock has a large effect.

of literature built on this work and, for the most part, focuses on the predictive power of financial variables, such as the term spread.[2]

A second strand of the literature explores the link between oil price changes and economic activity. While this literature is too voluminous for a complete discussion here, the central notion is that increases in oil prices have a detrimental effect on economic activity.[3] Over time, the literature has refined both its notion of an oil shock and the magnitude of its effects. Hamilton (1983), for example, considered all oil price innovations symmetrically; later, Mork (1989) found that only increases in oil prices had any effect. Recently, some studies [e.g., Hooker (1996) and Blanchard and Galí (2008)] have argued that oil's effect on the macroeconomy has weakened in the U.S., particularly during the Great Moderation period. This stylized fact has led some [e.g., Lee, Ni, and Ratti (1995) and Hamilton (1996, 2003, 2005)] to explore whether oil's effect is nonlinear as well as asymmetric. Many recent papers have adopted the net-oil-price-increase (NOPI) specification suggested in Hamilton (2003). He defines an oil shock as the magnitude of the increase in oil prices above the maximum price over the last three years.[4]

The intersection of the literature is the focus of this paper. While the overlap is sparse, a few authors have investigated the effect of oil prices on the timing of recessions. In particular, Raymond and Rich (1997) and Clements and Krolzig (2002) are among the more notable papers in the spirit of ours. Raymond and Rich (1997) used a Markov-switching model with time-varying transition probabilities [e.g., Filardo (1994) and Filardo and Gordon (1998)] to predict recessions. In their model, NOPI shocks may affect both the likelihood of recession (through the transition probability) and the level of output. Using in-sample diagnostics, they find that oil price increases affect the level of output but not the transition

[2] See Kauppi and Saikkonen (2008) and Wheelock and Wohar (2009) and the references therein.

[3] For a survey of the recent literature on oil price shocks, see Kliesen (2008) and the references therein. See also Kilian (2008a).

[4] Some have argued against asymmetric effects. Kilian and Vigfusson (2009) showed that assuming asymmetric effects in a VAR can lead to improper inference. Kilian (2009) contended that the oil price fluctuations are not, as often believed, exogenous shocks. He argued that shocks to oil supply and demand must be separately identified.

probabilities between the high- and low-growth phases. Clements and Krolzig (2002) used a three-state Markov-switching VAR to test whether oil prices can explain business cycle asymmetries. Their findings largely concur with Raymond and Rich (1997), that oil prices cannot adequately explain business cycle asymmetries.

In this paper, we reconsider an oil price shock's effect on the probability of recession in the U.S. and extend the analysis internationally.[5] Our model shares a number of similarities with the time-varying transition probability model of Raymond and Rich (1997). We adopt a Markov-switching model in which the transition probabilities vary with both net oil shocks and the term spread. Our contribution beyond Raymond and Rich (1997) – in addition to the international analysis – is the use of out-of-sample techniques for model evaluation. The salient result is that, for most countries, oil shocks are useful for assessing the probability of recessions out-of-sample. These results are not necessarily at odds with those of Raymond and Rich (1997) – instead, they suggest that oil prices may be less important for determining the historical incidence of recession and more useful for forecasting future recessions. In our model, these shocks must be both large and persistent in order to substantially influence the business cycle regime. We also find, for an admittedly small sample of countries, some variation in the strength of oil's effect.

The balance of the paper is laid out as follows: Section 2 describes the model used to estimate the recession dates jointly with the effect of oil shocks. Section 3 outlines the Gibbs sampler used to estimate the model. Section 4 presents the results for the U.S., comparing the estimated turning points with those supplied by the NBER. Section 5 presents the results for a few OECD countries. Section 6 concludes.

[5] A number of studies have sought to find international evidence of an effect of oil shocks on output. In early work, Burbidge and Harrison (1984) found that industrial production of only two out of the five countries studied (the U.S. and Japan) was negatively affected by oil shocks. In recent work, Kilian (2008b) studied the contributions of individual oil supply shocks to output growth for which he used exogenous shocks to oil production (with dates determined by major political events) and GDP growth rates for the G7 countries. Kilian found a significant negative effect on GDP growth within the first two years after an oil supply shock for most countries. Mork, Olsen, and Mysen (1994) found significant asymmetric effects of oil shocks for a number of countries. Similarly, Cuñado and Pérez de Gracia (2003) found that, for most of the European countries studied, an increase in oil prices resulted in a significant decrease in IP growth in the short run, but decreases in oil prices did not have the opposite effect.

2 Empirical Model

Many empirical models investigating oil's effects assume a linear relationship between oil prices and economic activity. Even in models in which oil shocks are defined as nonlinear permutations of oil prices [e.g., Hamilton (1996, 2003, 2005)], the shocks typically are assumed to affect economic activity linearly. Suppose, instead, that we believe large oil shocks cause sharp recessions, substantial and persistent declines in the growth rate of economic activity.

For the U.S., recessions can be thought of as the periods between the peaks and troughs defined by the NBER Business Cycle Data Committee. These dates provide the binary left-hand-side data that would be used to run, say, a probit-style regression. Taking the NBER dates as given, we can define a latent business cycle indicator S_t, where $S_t = 1$ if the economy is in recession and $S_t = 0$ if the economy is in expansion. The probit model computes the probability that the economy is in recession given a set of macroeconomic and financial covariates, X_t, which may include oil shocks, O_t, and the term spread, R_t, among other variables:

$$\Pr\left[S_t = 1 | X_t\right] = \Phi\left(X_t' \beta\right), \tag{1}$$

where β is a vector of coefficients and the link function, $\Phi(.)$, is the standard normal CDF.

In order to expand this type of analysis beyond the U.S., we require a set of binary recession indicators of the form provided us by the NBER. Unfortunately, no such data are available for both the cross-section of countries and the length of time series we desire.[6] As an alternative, we can estimate the underlying state variable by assuming that it follows a first-order Markov process with time-varying transition probabilities [see Goldfeld and Quandt (1973); Filardo and Gordon (1998)]. We could then assume the transition probabilities are

[6] The CEPR defines some dates for the Euro area. The Bank of England identifies years in which the UK experienced recessions but does not identify the quarter in which the turning point occurred. The OECD defines turning points for member countries using a leading indicator index. However, the OECD dates differ considerably from the NBER for the U.S. and from the Bank of England for the UK.

functions of the driving covariates, X_t, thereby unifying the spirit of the probit with the standard hidden Markov variable approach.

Formally, suppose the growth rate of GDP was given by

$$y_t = \mu_0 + \mu_1 S_t + \varepsilon_t, \tag{2}$$

where $S_t = \{0, 1\}$, $\varepsilon_t \sim N(0, \sigma^2)$, and $\mu_1 < 0$ prevents label switching. If we assumed that S_t followed a standard Markov process with constant transition probabilities, (2) represents the familiar switching model popularized by Hamilton (1989). This simple approach has been shown to identify turning points consistent with those defined by the NBER for U.S. data. On the other hand, suppose the latent variable, S_t, evolves according to a set of time-varying transition probabilities:

$$\begin{aligned} \Pr[S_t = 0 | S_{t-1} = 0, X_t] &= p(X_t), \\ \Pr[S_t = 1 | S_{t-1} = 1, X_t] &= q(X_t), \end{aligned} \tag{3}$$

which are (jointly) determined by the past state, S_{t-1}, and X_t. We can define a second (continuous) latent variable, S_t^*, as

$$S_t^* = \beta_0 + \beta_1 S_{t-1} + \beta_2 X_t + u_t, \tag{4}$$

where $u_t \sim N(0, 1)$ provides the normalization identifying the latent slope coefficients and we restrict $S_t^* > 0$ iff $S_t = 1$. The result is essentially a threshold model in which the binary latent business cycle indicator, S_t, determines the average growth rate of (say) GDP for two regimes, expansion and recession. The timing of those regimes is, in turn, influenced by the vector of covariates, X_t, through S_t^*. In this case, we are interested in whether increases in oil prices induce recessions and, thus, allow for the the covariates in X_t to affect the transition probabilities asymmetrically. That is, we can write

$$S_t^* = \beta_0 + \beta_1 S_{t-1} + \beta_2 (1 - S_{t-1}) X_t + u_t, \tag{5}$$

where, now, X_t has influence only when the economy is in expansion.[7]

3 Estimation

The model outlined in the preceding section can be estimated by the Gibbs sampler [e.g., Gelfand and Smith (1990) and Casella and George (1992)]. The Gibbs sampler makes iterative draws from the conditional distributions of subsets of the parameters to approximate the posterior distribution of the joint parameter vector, Θ. The sampler is constructed with four blocks: the business cycle average growth rates, $\boldsymbol{\mu} = (\mu_0, \mu_1)$; the innovation variance, σ^2; the transition probability parameters, $(\beta_0, \beta_1, \beta_2)$; and the vectors of latent variables, $\mathbf{S}_T$ and $\mathbf{S}_T^*$.

3.1 Priors

The sampler requires a set of priors for estimation. We adopt a normal-Gamma prior for parameters in (2): that is, $\boldsymbol{\mu} = [\mu_0, \mu_1]'$ has a prior distribution of the form $\boldsymbol{\mu} \sim \mathbf{1}_{[\mu_1<0]} N(\mathbf{m}_0, \sigma^2 \mathbf{M}_0)$, where $\mathbf{1}_{[\mu_1<0]}$ is an indicator function and denotes a truncation on the recession average growth rate that prevents label switching. The prior distribution for σ^2 is inverse-Gamma, $\sigma^{-2} \sim \boldsymbol{\Gamma}(\nu_0, d_0)$. The prior for the latent variable equation parameters, $\boldsymbol{\beta} = (\beta_0, \beta_1, \beta_2)$, is normal, $\boldsymbol{\beta} \sim N(\mathbf{b}_0, \mathbf{B}_0)$. The prior is parameterized such that $\mathbf{m}_0 = [1, -2]'$, $\mathbf{M}_0 = \mathbf{I}_2$, $\nu_0 = 1$, $d_0 = 1$, $\mathbf{b}_0 = \mathbf{0}_{p+2}$, and $\mathbf{B}_0 = diag(0.1, 0.1, 0.04 \times \mathbf{I}_p)$.

[7] We, thus, preclude *ex ante* the possibility that oil shocks influence recoveries. This interpretation is consistent with evidence from Mork (1989) and others.

3.2 The Sampler

After initializing the sampler, we draw μ_0, μ_1 conditional on $\mathbf{Y}$, $\mathbf{X}$, and $\Theta_{-\boldsymbol{\mu}}$, where $\Theta_{-\boldsymbol{\mu}}$ is the full parameter vector with $\boldsymbol{\mu}$ excluded. Define $\widetilde{\mathbf{X}} = [\mathbf{1}, \mathbf{S}_T]$. Then, $\boldsymbol{\mu} = [\mu_0, \mu_1]'$ is drawn from

$$\boldsymbol{\mu}|\mathbf{S}_T, \sigma_i \sim \mathbf{1}_{[\mu_1<0]} N\left(\mathbf{m}, \sigma^2\mathbf{M}\right),$$

where $\mathbf{M} = \left(\sigma\mathbf{M}_0 + \widetilde{\mathbf{X}}'\widetilde{\mathbf{X}}\right)^{-1}$ and $\mathbf{m} = \mathbf{M}\left(\sigma^{-2}\mathbf{M}_0^{-1}\mathbf{m}_0 + \widetilde{\mathbf{X}}'\mathbf{Y}\right)$.

Next, we draw σ^2 conditional on $\mathbf{Y}$, $\mathbf{X}$, and $\Theta_{-\sigma^2}$. The innovation variance can be drawn from the inverse-gamma posterior

$$\sigma^{-2}|\mathbf{Y}, \mathbf{X}, \Theta_{-\sigma^2} \sim \boldsymbol{\Gamma}\left(\frac{\nu_0 + T}{2}, \frac{d_0 + \widehat{\varepsilon}'\widehat{\varepsilon}}{2}\right),$$

where $\widehat{\varepsilon}_t = y_t - \mu_0 - \mu_1 S_t$ and T is the sample size.

We next draw $\mathbf{S}_T = \{S_1, ..., S_T\}$ recursively from

$$\Pr\left(S_t = 1|\mathbf{Y}, \mathbf{X}, S_{t+1}, S_{t-1}\right) \propto \Pr\left[S_t|S_{t-1}, X_{t-1}\right] \Pr\left[S_{t+1}|S_t, X_t\right] f\left(y_t|S_t\right),$$

where $\Pr\left[S_t|S_{t-1}, X_{t-1}\right]$ and $\Pr\left[S_{t+1}|S_t, X_t\right]$ are the values of the time-varying transition probabilities defined by (3) and (4). Conditional on S_t and $\boldsymbol{\beta}$, the latent variables, $\mathbf{S}_T^*$, can be drawn from truncated normal distributions.

Finally, the latent state parameters, $\boldsymbol{\beta}$, are drawn from

$$\boldsymbol{\beta} \sim N\left(\mathbf{b}, \mathbf{B}\right),$$

where $\mathbf{B} = (\mathbf{B}_0 + \mathbf{W}_T'\mathbf{W}_T)^{-1}$, $\mathbf{b} = \mathbf{B}\left(\mathbf{B}_0^{-1}\mathbf{b}_0 + \mathbf{W}_T'\mathbf{S}_T^*\right)$, $\mathbf{S}_T^* = [S_1, ..., S_T]'$, and $\mathbf{W}_T$ is a $T \times (2+p)$ matrix with rows $[1, S_{t-1}, (1 - S_{t-1}) X_t]$. We repeat these draws $10,000$ times and discard the first $5,000$ draws. The remaining draws form the joint posterior for the full parameter vector.

3.3 Data

Our economic indicator is the seasonally-adjusted annualized growth rate of real GDP obtained at a quarterly frequency. Data availability limits our cross-section of countries and overall sample period. Table 1 lists the set of countries in our sample, information on the data series, and the data sources. Our baseline model includes the term spread – in most cases defined as the 10-year and 3-month Treasury equivalents for each country.

Consistent with many recent studies, we use the net oil price increase (NOPI) over a three-year period as our measure of innovations to the oil prices. The net oil price is computed as the maximum of zero and the log of the ratio of the current oil price to the maximum of the last three years. For the U.S., previous studies have used two different spot price measures – West Texas Intermediate (WTI) and the producer price index for oil – with, for the most part, similar results. Results shown for the U.S. in the next section are for WTI; however, alternate oil series produced similar results. For the cross-country analysis, we used the NOPI computed from the spot price of world crude oil series supplied by the International Monetary Fund (IMF). For comparison, we reproduce the U.S. results using the IMF oil prices.

4 U.S. Results

The model described in the preceding section is designed to simultaneously identify switches in the business cycle regime and assess the influence of variables in X_t on transitions between these regimes. The output of the sampler is a time series of posterior recession probabilities – i.e., $\Pr[S_t = 1|\mathbf{Y}, \mathbf{X}]$ – and the posterior distributions of the model parameters – i.e., $p(\Theta|\mathbf{Y}, \mathbf{X})$.[8] For the U.S., we estimate the model for a number of combinations of oil shock and term spread lags in the latent variable equation.

[8] Obviously, the effect of idiosyncratic shocks that may cause recessions (for example, the bursting of housing or asset price bubbles) but are not included in X_t cannot be assessed. The model can be easily expanded to include these variables. We leave that for future research.

Model choice is made utilizing k-fold cross validation [see, for example, Gelfand, Dey, and Chang (1992)], a method similar to out-of-sample forecast evaluation but using the full set of data. Cross validation partitions the full set of data into k subsamples. For each subsample, we estimate the model with the ith subsample omitted; we, then, use the ith subsample to compute a prediction loss (say, an entropy measure). The sum of these losses over each of the k subsamples yields a score for the specification with the lowest score being preferred. Cross validation essentially asks how well the ith subsample could be fit using the model estimated with the remaining data. These methods have been adopted (and modified) for similar models in econometrics to test specifications where Bayes factors are computationally taxing [see Geweke and Keane (2007)]. For the U.S., we computed the cross validation score for all specifications between 0 (contemporaneous) and 12 lags each of the WTI NOPI and the term spread.

Table 2 reports the cross validation scores for a subset of the model specifications tested. In almost all cases, including the contemporaneous NOPI and the contemporaneous term spread improves the cross validation score. In all cases, models with time-varying transition probabilities beat the constant probability Markov model. We find the two best specifications include 10 lags of oil and either 1 (second best) or 7 (best) lags of the term spread.[9] Previous work relating business cycle turning points to the term spread has generally called for more lags of the term spread.

One way of assessing the model's veracity is to compare the posterior recession probabilities to the recessions identified by the NBER Business Cycle Dating Committee. Figure 1 plots the NBER dates along with the time series of posterior probabilities for the model including ten lags of the WTI NOPI series, seven lags of the term spread, contemporaneous oil, and contemporaneous term spread. We include the oil shocks as defined by Hamilton (2003) for reference. For six of the ten NBER-identified recessions during the sample period, the model identifies the business cycle peak – give or take a few quarters – with at least a

[9] The cross validation score for the best specification was 291.89, and the score for the constant probability Markov model was about twice as large — 599.86.

50-percent posterior probability. In two of the remaining four cases (1990-91 and 2001), the probability of a recession rises around the NBER peak to between 20 percent and 40 percent. The latter case was preceded by several rather small oil shocks while the former case, on the other hand, was preceded by a rather large shock during the NBER peak quarter. One potential explanation for this result is that an oil-triggered increase in the probability of a recession occurs only when energy prices rise by a rather large amount over a number of quarters.[10] A second possibility is that the model treats all recessions identically; the recessions occurring in 1953-54, 1960-61, 1990-91, and 2001 were shallow compared to others in the sample period.

In addition to "missing" some recessionary episodes, the model identifies a period in the mid-1960s that is not included in the NBER dates as recessionary. This apparently false recession was characterized by a number of small oil price shocks, an extended inversion of the term spread, and very low but positive GDP growth. While the NBER did not declare this period a recession, the OECD, using a different methodology, did call 1966:III a peak.

A third characteristic of Figure 1 is that the model, in some cases, identifies business cycle turning points at different times than the NBER. This is not surprising considering the NBER uses a different set of business cycle indicators and emphasizes more series than GDP growth alone. Nevertheless, we believe the model still performs qualitatively well and can, at the very least, be used to assess the effect of net oil price shocks on recession probabilities.

Table 3 provides summary statistics for the model parameters for (2).[11] Coefficients on the oil and term spread lags in the latent variable equation have the expected sign. As we previously indicated, the latent variable equation, (5), has a similar interpretation as the standard probit model – that is, we can use the parameter values to determine the effect of a shock to X_t on the transition probability. Because X_t is a vector that includes lagged values,

[10] The 1953-54 NBER recession was preceded by a one-quarter NOPI shock of less than 10 percent. The 1960-61 NBER recession was not preceded by an oil shock. In both cases, the model found almost no increase in the posterior probability of a recession.

[11] In the interest of brevity, we do not report the lag coefficients for oil and the term spread. The full set of results are available upon request.

an increase in X at time t will have effects on the transition probabilities for a number of periods. This makes disentangling the effect of a one-time oil shock more difficult.

We can, however, compute a counterfactual increase in the probability of going into a recession p periods ahead caused by a one-percentage-point increase in net oil prices at time t. We construct this counterfactual probability assuming that the economy has been in expansion ($S_t = 0$) up through period $t + p - 1$, that the term spread is constant for all time, and that no other net oil shocks occur. This experiment is conducted by setting all of the oil shocks except the one at time t equal to zero and setting the term spread equal to a constant, r, for all time. This yields a change in the recession probability for various (constant) levels of the term spread, r:

$$\left.\frac{\partial \Pr\left[S_{t+p} = 1 | S_{t+p-1} = ... = S_{t-1} = 0, X_{t+p}\right]}{\partial O_t}\right|_{\mathbf{R}_{t+p} = r\mathbf{1}_{t+p}; \mathbf{O}_{-t} = \mathbf{0}_{t+p-1}} = \phi\left(\beta_0 + r\beta_2^R \mathbf{1}_q + \beta_{2p}^{oil} O_t\right)\beta_{2p}^{oil}, \tag{6}$$

where O_t is the period$-t$ oil shock, $\mathbf{O}_{-t}$ is the vector of oil shocks for all periods except t, $\mathbf{R}_{t+p} = [R_{t+p}, ..., R_1]$, β_2^R is the $(1 \times q)$ subvector of β associated with q lags of the term spread, $\mathbf{1}_m$ is an $(m \times 1)$ vector of ones, $\mathbf{0}_m$ is an $(m \times 1)$ vector of zeros, β_{2p}^{oil} is the coefficient on the pth lag of the oil shock, and $\phi(.)$ is the normal pdf. Based on (6), we can plot the (counterfactual) evolution of the transition probability's (cumulative) response to an oil shock at time t, assuming that no recession has occured yet. Figure 2 shows this evolution for $p = 0, 1, ..., 10$, calibrated to the empirical results using the same specification as in Figure 1 (the final column of Table 3) for various term speads. Note that the timing of the turning point need not be coincident with the incidence of the oil shock. The transition probability rises from the time of the oil shock through period 10 before slowly leveling off.[12] During the sample period, the average shock to the price of WTI oil is about 13 percentage

[12] It is important to keep in mind that Figure 3 represents the transition probability conditional on the economy staying in expansion throughout. Once the economy switches to a recession, the transition probability evolves differently.

points over the previous 3-year maximum price. Thus, for a zero term spread, the baseline transition probability from expansion to recession would rise nearly 50 percentage points after one year, nearly 90 percentage points after two years, and over 100 percentage points after 10 quarters, all else equal.

At a glance, these results appear to stand in contrast to Raymond and Rich's (1997) finding that oil price shocks do not influence the timing of switches in business cycle regimes. However, quick consideration of their Table 2 shows that their results are not substantially different from ours.[13] Raymond and Rich's interpretation that oil has no effect on the timing of the switches is based on likelihood ratio scores comparing the time-varying transition probability model to the constant transition probability model. In addition, their conclusions are based on the *in-sample* results. We find that, for the various specifications in Table 3, estimation using the full sample yields, essentially, the same posterior regime probabilities. That is, in-sample, there appears to be little information gained by adding variables to the latent equation, (4). Using the "out-of-sample" metric, however, we find that any model with oil prices dominates the constant transition probability model, typically halving the validation score. This suggests that oil prices may not be more informative than, say, GDP data alone for determining historical business cycles but may be important for forecasting future turning points.

5 International Evidence

One of the advantages of *estimating* business cycle turning points is that the model can be extended to countries for which business cycle data is unavailable. In this section, we describe the results for some OECD countries with sufficiently long time samples of both real GDP growth and a term spread: Australia, Canada, France, Japan, Norway, and the UK. We use the IMF oil series for our international analysis and reestimate the U.S. turning

[13] Raymond and Rich's preferred TVTP specification includes the third and fourth quarterly lag of oil prices. More recent innovations in oil prices are excluded. They also do not include the term spread variable.

points for comparison.

As with the U.S., we can assess the qualitative performance of the model by examining the posterior recession probabilities. Figure 3 shows these for the countries in question. As expected, the use of a common oil shock series leads to a number of coincident recessions across countries. In particular, the recent downturn beginning in 2007 was indeed global. However, there is also a significant amount of cross-country variation in the posterior regime probabilities. For example, France appears to have experienced a prolonged recession in the early 1990s, a period which was called a recession in the U.S. by the NBER Business Cycle Dating Committee but we showed previously exhibits only a small increase in the recession probability (around 40 percent). All countries with available data experienced at least an 80 percent probability of a peak around 1974, which coincided with the largest oil shock in our sample — about 84 percentage points. Except for Japan, they all had at least one quarter in the early 2000s when the probability of a peak was greater than 60 percent; there were shocks to world crude oil four quarters in a row — 1999:IV-2000:III — between 6 and 10 percentage points. The probability of a peak remained less than 15 percent in Japan from 1992:II through 2008:I.

Table 4 provides summary statistics for some of the parameter values when estimating the model with data from various countries.[14] Note that the preferred specification, while typically including at least nine lags of the NOPI, differs across countries. Except for Australia, Canada, and the UK, the preferred model also includes contemporaneous oil shocks and term spreads. The ratio of the preferred specification's validation score to that of the constant transition probability model ranges from 0.33 for Japan to 0.74 for Norway, a strong indication of oil's predictive power for the countries in our sample. In addition to the differences in specification, the magnitude of oil's effect varies across countries.

Some papers suggest differentiating between oil-importing and oil-exporting countries. Peersman and Van Robays (2009) studied three types of oil shocks—oil supply shocks, oil de-

[14] The full set of results is again available upon request.

mand shocks induced by increased world output, and other oil demand (oil-specific) shocks—on GDP. They found that oil supply shocks lead to a permanent decrease in GDP for net energy importers and a permanent increase in GDP for net energy exporters. With the oil demand shocks, for all countries they found a temporary increase in GDP when the shock was GDP-driven but a temporary decrease in GDP for oil-specific shocks. Jiménez-Rodríguez and Sánchez (2005) found that the oil-importing countries in their study (except Japan) experience a decline in GDP growth in response to an oil price increase. The oil-exporting countries (Norway and the UK) experience an initial increase in GDP growth followed by a decline after a positive oil shock. However, the overall impact of an increase in oil prices on Norway's GDP is positive while the overall impact on the UK's GDP is negative.

Figure 4 plots the total increase in the transition probability induced by a one-percentage-point NOPI shock against each country's net oil exports per capita. Norway – the country most responsive to oil shocks – is far and away the most substantial exporter of oil in this group. It also has some of the noisiest quarterly-level data. Moreover, some have argued that exporters of large quantities of oil should be treated differently than other countries. In contrast, the other two net exporters – Canada and the UK – are the least responsive to oil shocks. While by no means conclusive, the general tendency is for lower oil exports (alternatively, higher oil imports) per capita to increase oil's effect on a country's recession probability.

As in the previous section, we used the preferred specification in Table 4 to compute counterfactual increases in transition probabilities. Figure 5 shows the cumulative effects of a one-percentage-point NOPI shock for each country. For zero term spread, the overall effects are fairly similar (between 4 and 6 percentage points) for the majority of countries, but the effects for Japan, Canada, and Norway stand in contrast. The overall effects for Japan are roughly 1.5 to 2 times larger than for the other countries, while the overall effects for Canada are roughly one-half to two-thirds smaller than for the other countries. For Norway, the effects are relatively small at zero but very large when the yield curve is inverted. For

Japan and Canada, an average shock to the world crude oil price at a zero term spread increases the probability of recession by 115 and 28 percentage points, respectively, over ten quarters.[15] For the remaining countries, the average NOPI shock increases the baseline probability of a peak between 64 and 82 percentage points.

6 Conclusion

Many statistical models of business cycles characterize the expansion and recession phases as constant transition probability Markov processes. While these methods have proven useful for identifying historical business cycle phases, recent papers aimed at predicting recessions out-of-sample find it beneficial to model the transition probabilities as functions of some exogenous drivers. Many of these recent studies have shown that the term spread – a measure of expectations – is a useful predictor of future turning points. In contrast to some work, we show that net oil prices also have predictive content for determining turning points.

We find that, for the U.S., the specification that predicts the business cycle best "out-of-sample" includes both lags of oil prices and lags of the term spread. This result extends to the small sample of OECD countries we tested, although the overall effect of oil price shocks is varied. Our findings are based on the use of a quasi–out-of-sample measure, k-fold cross validation, which considers how the estimated model fits subsamples of the data.

If, indeed, there exist large asymmetries in business cycle phases, policymakers and financial markets have incentives to predict upcoming turning points. Our results imply that oil prices do have some predictive ability for forecasting recessions and that oil price shocks — in addition to the term spread — can be used to date current turning points.

[15] Because Japan's and Norway's sample periods are shorter than the other countries, the average shock to the IMF NOPI is slightly smaller for the two (13 percent versus 15 percent).

References

[1] Blanchard, Olivier J. and Jordi Galí (2008) The Macroeconomic Effects of Oil Price Shocks: Why Are the 2000s So Different from the 1970s? Working Paper.

[2] Burbidge, John and Alan Harrison (1984) Testing for the Effects of Oil-Price Rises Using Vector Autoregressions. *International Economic Review* 25(2), 459-484.

[3] Casella, George and Edward I. George (1992) Explaining the Gibbs Sampler. *American Statistician* 46(3), 167-174.

[4] Clements, Michael P. and Hans-Martin Krolzig (2002) Can Oil Shocks Explain Asymmetries in the US Business Cycle? *Empirical Economics* 27(2), 185-204.

[5] Cuñado, Juncal and Fernando Pérez de Gracia (2003) Do Oil Price Shocks Matter? Evidence for some European Countries. *Energy Economics* 25(2), 137-154.

[6] Estrella, Arturo and Frederic S. Mishkin (1998) Predicting U.S. Recessions: Financial Variables as Leading Indicators. *Review of Economics and Statistics* 80(1), 45-61.

[7] Filardo, Andrew J. (1994) Business-Cycle Phases and Their Transitional Dynamics. *Journal of Business and Economic Statistics* 12(3), 299-308.

[8] Filardo, Andrew J. and Stephen F. Gordon (1998) Business Cycle Durations. *Journal of Econometrics* 85(1), 99-123.

[9] Gelfand, Alan E., Dipak K. Dey, and Hong Chang (1992) Model Determination Using Predictive Distributions with Implementation via Sampling-Based Methods. In J. M. Bernardo et al. (eds.), *Bayesian Statistics 4: Proceedings of the Fourth Valencia International Meeting*, pp. 147-159. Oxford: Oxford University Press.

[10] Gelfand, Alan E. and Adrian F. M. Smith (1990) Sampling-Based Approaches to Calculating Marginal Densities. *Journal of the American Statistical Association* 85(410), 398-409.

[11] Geweke, John and Michael Keane (2007) Smoothly Mixing Regressions. *Journal of Econometrics* 138(1), 252-290.

[12] Goldfeld, Stephen M. and Richard E. Quandt (1973) A Markov Model for Switching Regressions. *Journal of Econometrics* 1(1), 3-15.

[13] Hamilton, James D. (1983) Oil and the Macroeconomy since World War II. *Journal of Political Economy* 91(2), 228-248.

[14] Hamilton, James D. (1989) A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle. *Econometrica* 57(2), 357-384.

[15] Hamilton, James D. (1996) This Is What Happened to the Oil Price–Macroeconomy Relationship. *Journal of Monetary Economics* 38(2), 215-220.

[16] Hamilton, James D. (2003) What Is an Oil Shock? *Journal of Econometrics* 113(2), 363-398.

[17] Hamilton, James D. (2005) Oil and the Macroeconomy. Prepared for: *Palgrave Dictionary of Economics.*

[18] Hooker, Mark A. (1996) What Happened to the Oil Price–Macroeconomy Relationship? *Journal of Monetary Economics* 38(2), 195-213.

[19] Jiménez-Rodríguez, Rebeca and Marcelo Sánchez (2005) Oil Price Shocks and Real GDP Growth: Empirical Evidence for Some OECD Countries. *Applied Economics* 37(2), 201-228.

[20] Kauppi, Heikki and Pentti Saikkonen (2008) Predicting U.S. Recessions with Dynamic Binary Response Models. *Review of Economics and Statistics* 90(4), 777-791.

[21] Kilian, Lutz (2008a) Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy? *Review of Economics and Statistics* 90(2), 216-240.

[22] Kilian, Lutz (2008b) A Comparison of the Effects of Exogenous Oil Supply Shocks on Output and Inflation in the G7 Countries. *Journal of the European Economic Association* 6(1), 78-121.

[23] Kilian, Lutz (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.

[24] Kilian, Lutz and Robert J. Vigfusson (2009) Pitfalls in Estimating Asymmetric Effects of Energy Price Shocks. Working Paper.

[25] Kliesen, Kevin L. (2008) Oil and the U.S. Macroeconomy: An Update and a Simple Forecasting Exercise. Federal Reserve Bank of St. Louis *Review* 90(5), 505-516.

[26] Lee, Kiseok, Shawn Ni, and Ronald A. Ratti (1995) Oil Shocks and the Macroeconomy: The Role of Price Variability. *Energy Journal* 16(4), 39-56.

[27] Mork, Knut Anton (1989) Oil and the Macroeconomy When Prices Go Up and Down: An Extension of Hamilton's Results. *Journal of Political Economy* 97(3), 740-744.

[28] Mork, Knut Anton, Oystein Olsen, and Hans Terje Mysen (1994) Macroeconomic Responses to Oil Price Increases and Decreases in Seven OECD Countries. *Energy Journal* 15(4), 19-35.

[29] Peersman, Gert and Ine Van Robays (2009) Cross-Country Differences in the Effects of Oil Shocks. Working Paper.

[30] Raymond, Jennie E. and Robert W. Rich (1997) Oil and the Macroeconomy: A Markov State-Switching Approach. *Journal of Money, Credit, and Banking* 29(2), 193-213.

[31] Wheelock, David C. and Mark E. Wohar (2009) Can the Term Spread Predict Output Growth and Recessions? A Survey of the Literature. Federal Reserve Bank of St. Louis *Review* 91(5, Part 1), 419-440.

Table 1: Data

	GDP Series	Source	Term Spread Series	Source	Oil Price Series	Source	Sample Period
Australia	SA, Mil.Chn.05-06.A$	OECD	10-Yr Commonwealth Treasury Bond Yield (% p.a.)	OECD	World (US$/Bbl)	IMF	1973:I–2009:I
			90-Day Commercial Bill Rate (% p.a.)				
Canada	SAAR, Mil.Chn.02.C$	OECD	10-Yr Benchmark Bond Yield (% p.a.)	OECD	World (US$/Bbl)	IMF	1965:I–2009:I
			90-Day Prime Corporate Paper (EOP, % p.a.)				
France	SA, Bil.Chn.00.Euros	OECD	10-Yr Gov. Bond Yields (% p.a.)	OECD	World (US$/Bbl)	IMF	1973:I–2009:I
			3-Month PIBOR (%)				
Japan	SAAR, Bil.Chn.00.Yen	OECD	10-Yr Central Gov. Bond Yield (% p.a.)	OECD	World (US$/Bbl)	IMF	1992:I–2009:I
			3-Month C.D. {Gensaki Rate}(% p.a.)				
Norway	SA, Mil.Chn.05 Kroner	OECD	10-Yr Central Gov. Bond Yield (% p.a.)	OECD	World (US$/Bbl)	IMF	1988:I–2009:I
			3-Month Interbank Offer Rate: NIBOR (%)				
UK	SA, Mil.Chn.03.Pounds	OECD	10-Yr Central Gov. Bond Yield (% p.a.)	OECD	World (US$/Bbl)	IMF	1964:I–2009:I
			90-Day Treasury Bills Rate (% p.a.)				
U.S.	SAAR, Bil.Chn.05$	BEA	Long-Term Treasury Composite >10 Yrs (% p.a.)	Treasury	WTI ($/Barrel)	WSJ	1952:I–2009:II w/ WTI
			3-Month Treasury Bills, Secondary Market (% p.a.)	FRB	World (US$/Bbl)	IMF	1964:I–2009:II w/ World

NOTES: The world spot oil price is the average of UK Brent Light, Dubai Medium, and Alaska NS heavy. WTI is the spot price of West Texas Intermediate.

We used the log growth rate of real GDP and net oil price increase. Our term spread series were calculated as the long-term rate minus the short-term rate.

Table 2: U.S. Cross Validation Score Ratios								
With contemporaneous oil and term spread								
# of oil lags	# of term spread lags (horizontal)							
(vertical)	1	2	4	6	7	8	10	12
1	0.57	0.63	0.61	0.63	0.62	0.62	0.61	0.59
2	0.57	0.63	0.61	0.60	0.61	0.60	0.60	0.58
4	0.54	0.60	0.57	0.57	0.56	0.56	0.57	0.56
6	0.52	0.60	0.54	0.55	0.54	0.54	0.56	0.54
8	0.51	0.59	0.54	0.54	0.54	0.54	0.53	0.53
10	0.49	0.54	0.54	0.53	**0.49**	0.50	0.57	0.54
12	0.62	0.62	0.76	0.70	0.73	0.73	0.73	0.51
Without contemporaneous oil and term spread								
# of oil lags	# of term spread lags (horizontal)							
(vertical)	1	2	4	6	7	8	10	12
1	0.66	0.66	0.66	0.67	0.67	0.66	0.67	0.63
2	0.66	0.67	0.67	0.67	0.67	0.66	0.66	0.64
4	0.65	0.66	0.64	0.63	0.63	0.63	0.64	0.62
6	0.65	0.65	0.62	0.62	0.62	0.62	0.63	0.61
8	0.64	0.65	0.62	0.60	0.62	0.62	0.63	0.60
10	0.63	0.64	0.59	0.60	0.60	0.63	0.63	0.62
12	0.62	0.65	0.62	0.64	0.60	0.61	0.61	0.56

NOTES: The WTI NOPI was used. The score ratio is the ratio of the specification's validation score to that of the constant transition probability model ('0_0_N_N'), which was 600. The preferred specification includes 10 lags of the NOPI, 7 lags of the term spread, and contemporaneous oil and term spread (shown in bold).

Table 3: U.S. Results, Model Parameters

Model	'4_1_N_N'	'8_4_N_N'	'10_7_N_N'	'4_1_Y_Y'	'8_4_Y_Y'	'10_1_Y_Y'	'10_7_Y_Y'
μ_0	0.91 (0.76,1.07)	0.95 (0.81,1.11)	0.93 (0.78,1.09)	0.90 (0.75,1.05)	0.93 (0.78,1.08)	0.91 (0.76,1.07)	0.93 (0.78,1.09)
μ_1	−1.06 (−1.47,−0.77)	−1.03 (−1.29,−0.80)	−0.98 (−1.24,−0.77)	−0.91 (−1.14,−0.72)	−0.99 (−1.23,−0.78)	−0.89 (−1.08,−0.73)	−0.96 (−1.20,−0.76)
σ^2	0.80 (0.65,0.95)	0.77 (0.65,0.91)	0.80 (0.67,0.95)	0.85 (0.72,1.00)	0.80 (0.67,0.94)	0.85 (0.72,1.00)	0.80 (0.67,0.95)
β_0	−0.14 (−0.60,0.25)	0.66 (0.29,1.03)	0.84 (0.44,1.28)	0.11 (−0.27,0.43)	0.75 (0.37,1.13)	0.41 (0.08,0.75)	0.95 (0.51,1.36)
β_1	0.20 (−0.06,0.61)	0.08 (−0.06,0.23)	0.10 (−0.05,0.24)	0.07 (−0.09,0.29)	0.09 (−0.06,0.23)	0.05 (−0.09,0.20)	0.11 (−0.04,0.26)
score ratio	0.65	0.62	0.60	0.54	0.54	0.49	0.49

NOTES: The model is the number of oil lags, number of term spread lags, contemporaneous oil, and contemporaneous term spread. The 5 percent and 95 percent error bands are shown in parentheses. The table includes the second-best and our preferred specifications (last two columns). The score ratio is the ratio of the specification's validation score to that of the constant transition probability model ('0_0_N_N'), which was 600. The WTI NOPI series was used for all of these results.

Table 4: Country Results, Preferred Specifications

Country	Australia	Canada	France	Japan	Norway	UK	U.S.
Model	'9_7_N_N'	'10_5_N_N'	'10_2_Y_Y'	'9_1_Y_Y'	'9_11_Y_Y'	'6_4_N_N'	'7_4_Y_Y'
μ_0	0.50 (0.40,0.81)	0.90 (0.63,1.15)	0.67 (0.54,0.80)	0.29 (0.09,0.49)	0.51 (−0.07,1.08)	0.70 (0.50,0.90)	0.94 (0.78,1.09)
μ_1	−0.48 (−0.75,−0.39)	−0.83 (−1.04,−0.62)	−0.64 (−0.78,−0.51)	−2.79 (−3.75,−1.95)	−0.54 (−1.06,−0.02)	−0.68 (−0.87,−0.48)	−0.95 (−1.17,−0.76)
σ^2	0.95 (0.71,1.21)	0.73 (0.59,0.91)	0.32 (0.25,0.39)	0.58 (0.42,0.77)	1.75 (1.25,2.40)	0.88 (0.73,1.05)	0.65 (0.54,0.79)
β_0	1.07 (0.45,1.82)	1.58 (1.05,2.23)	1.01 (0.53,1.57)	−0.26 (−0.57,0.05)	3.70 (1.77,6.22)	0.67 (0.33,1.03)	0.88 (0.44,1.32)
β_1	0.33 (0.07,0.61)	0.30 (0.10,0.51)	0.19 (0.002,0.40)	0.30 (0.03,0.59)	1.60 (0.66,2.75)	0.11 (−0.05,0.28)	0.13 (−0.04,0.29)
score ratio	0.56	0.53	0.37	0.33	0.74	0.61	0.45

NOTES: The model is the number of oil lags, number of term spread lags, contemporaneous oil, and contemporaneous term spread. The 5 percent and 95 percent error bands are shown in parentheses. The score ratio is the ratio of the preferred specification's validation score to that of the constant transition probability model ('0_0_N_N').

The world NOPI series was used for all of these results.

Figure 1: U.S. Posterior Probablities, NBER Recessions and Oil Shocks

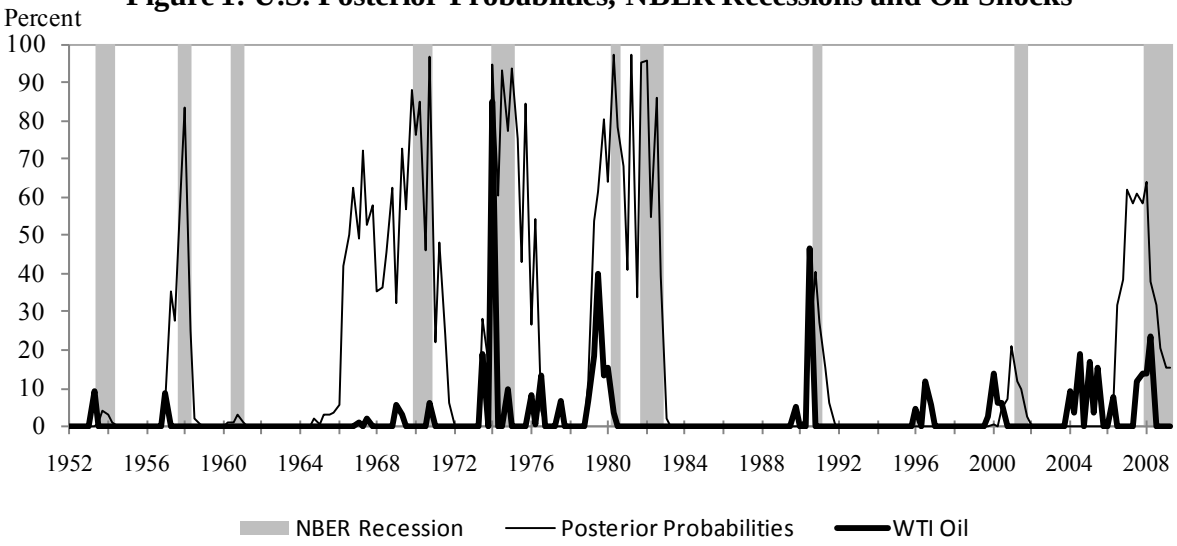

NOTES: The specification included 10 WTI NOPI lags, 7 term spread lags, contemporaneous oil prices, and contemporaneous term spread. The chart contains net oil price increases as defined by Hamilton (2003).

Figure 2: Increase in Posterior Probability in Response to Oil Shock

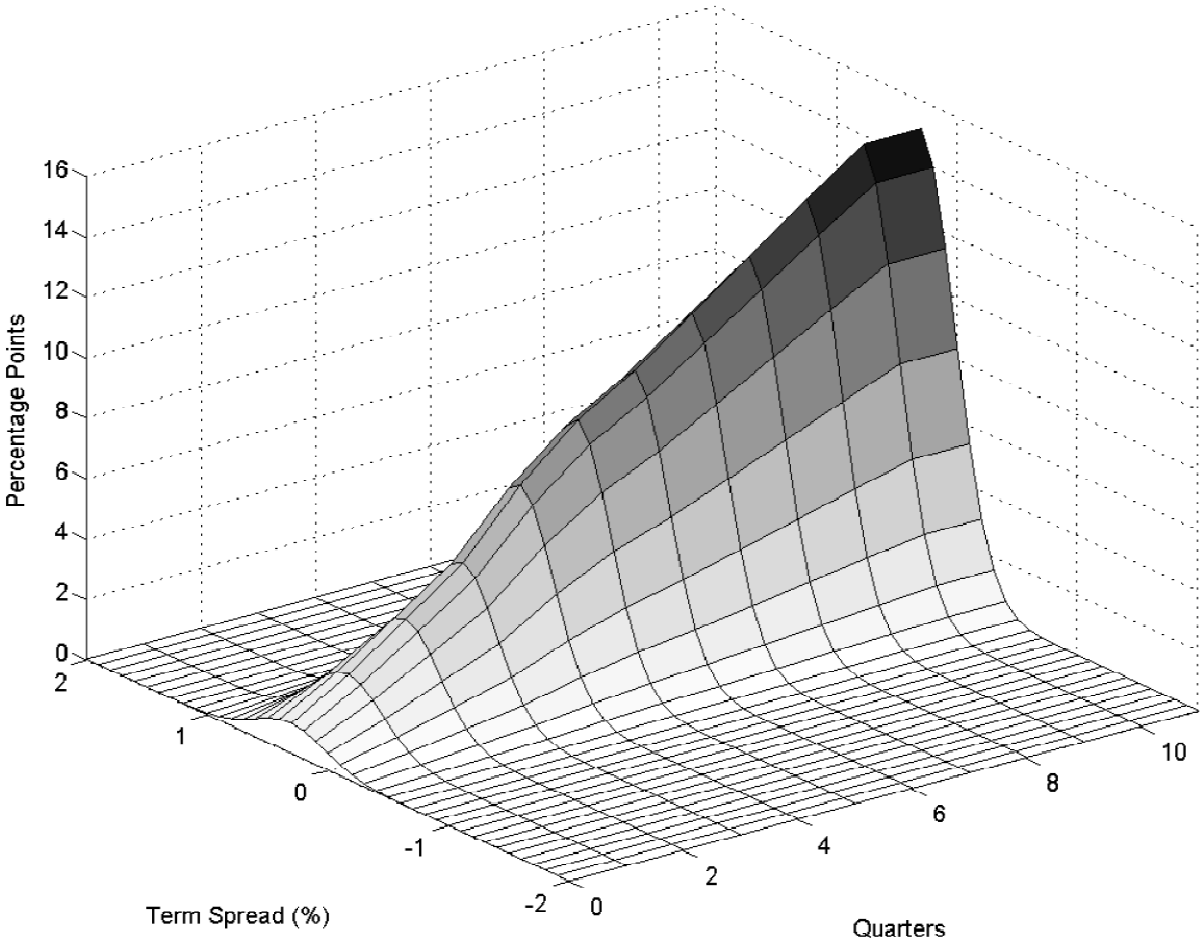

NOTES: These counterfactual responses are calibrated to the U.S. results using 10 WTI NOPI lags, 7 term spread lags and contemporaneous oil prices. The cumulative effects of a one-percentage-point NOPI shock are shown.

Figure 3: GDP vs. Posterior Recession Probabilities, Preferred Specifications

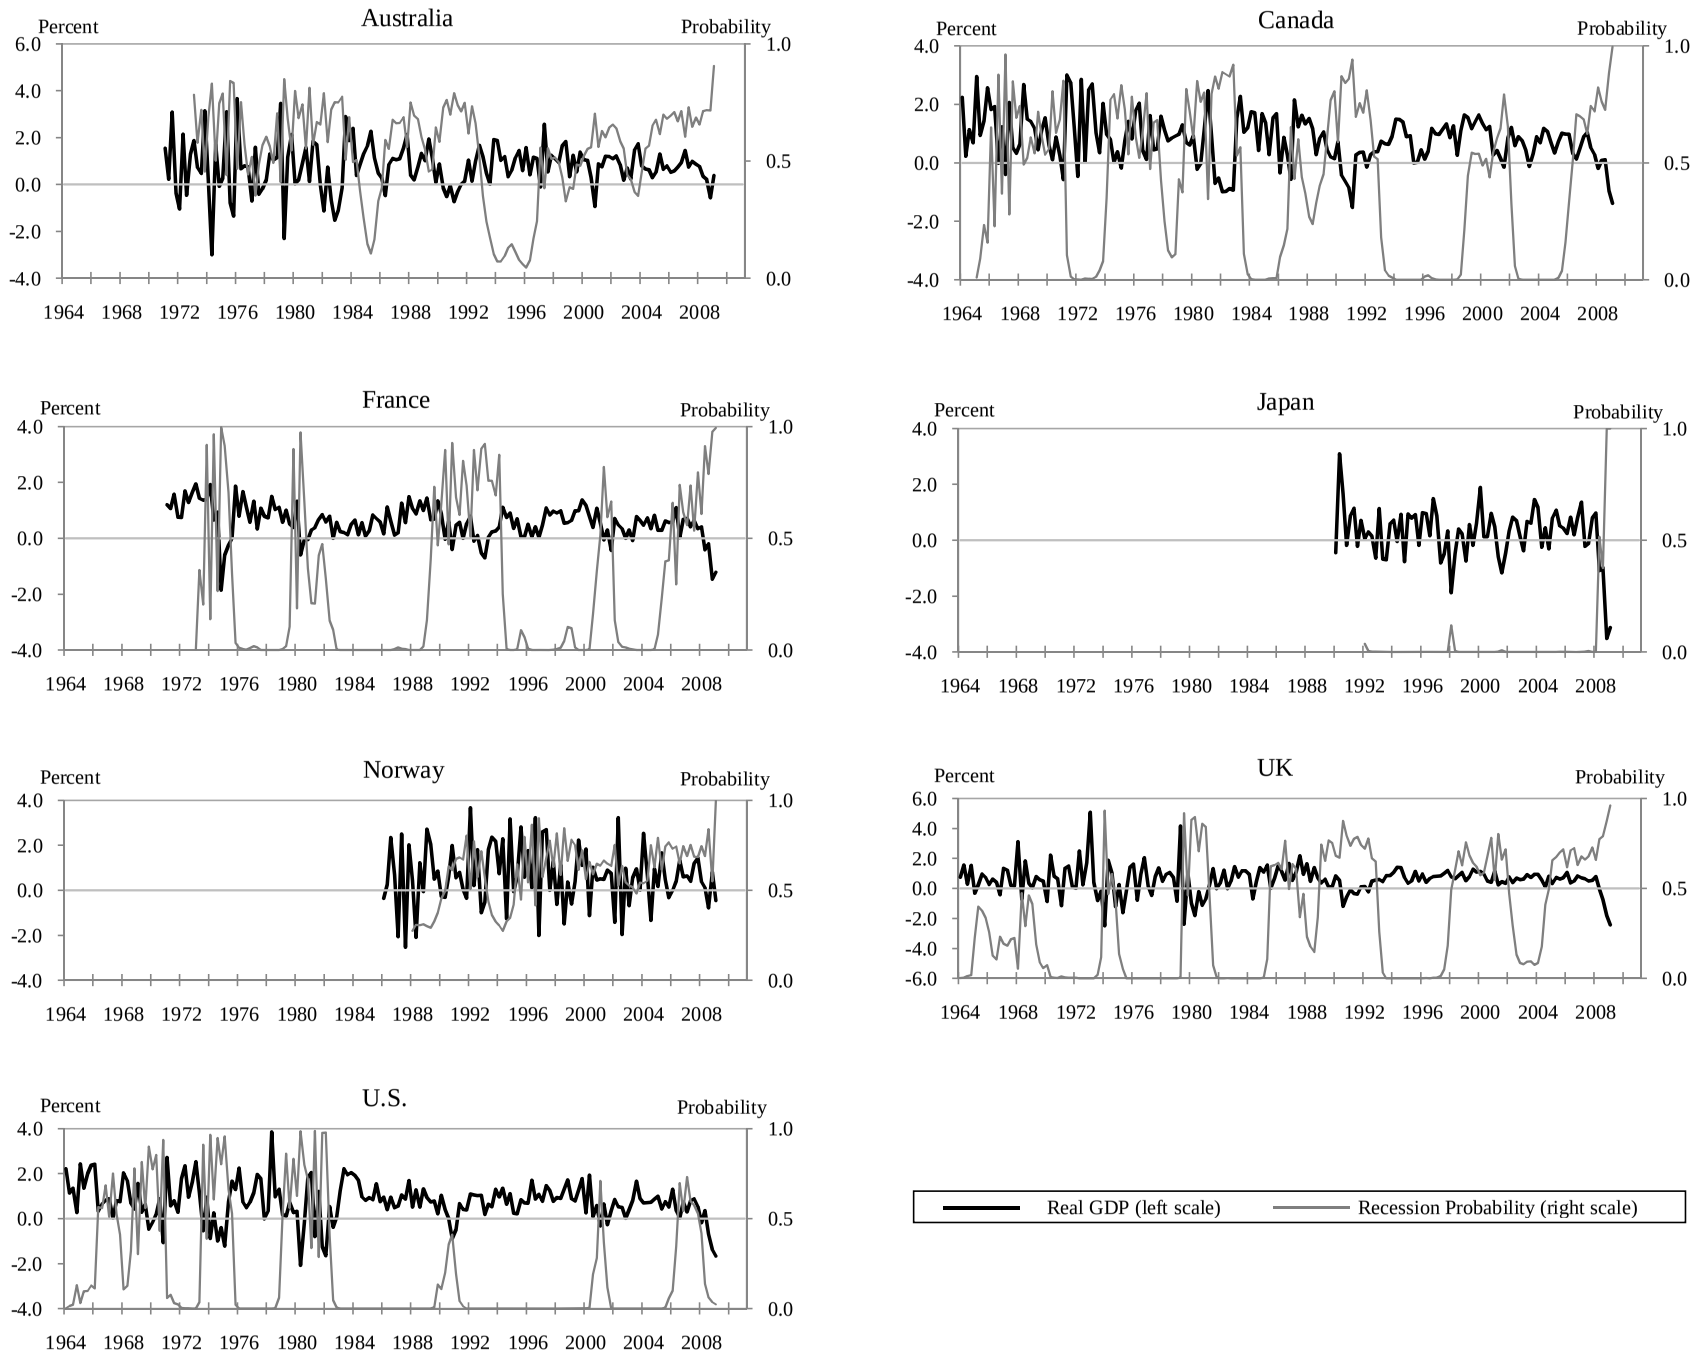

NOTES: For all countries, the world NOPI series and the log growth of real GDP were used.

Figure 4: Sum of Oil Coefficients vs. Net Oil Exports

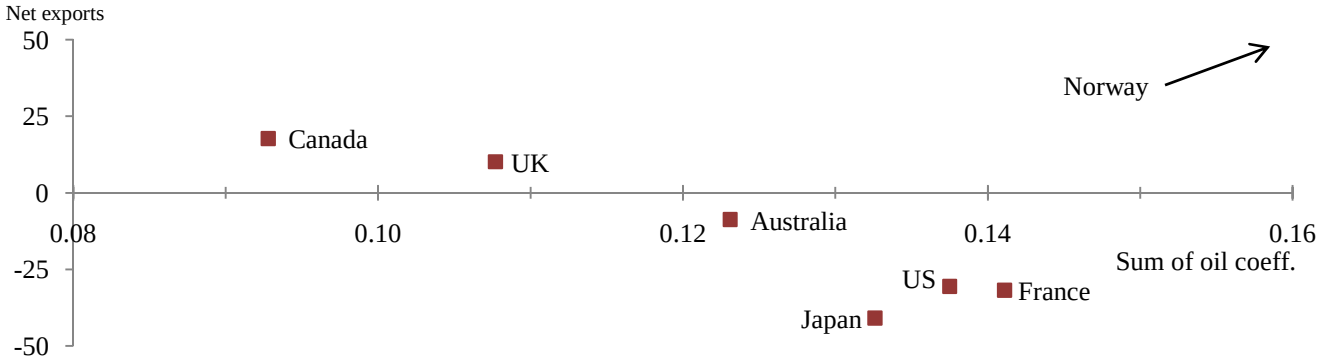

NOTES: For each country, the sum of oil coefficients is for the preferred specification in Table 5. Net oil exports were calculated from net petroleum exports in barrels/day (from the U.S. Energy Information Adminstration) and population (from the IMF). The series plotted is net exports per day per 1,000 people averaged over 1980-2008; negative numbers indicate net imports. For Norway, the sum of oil coefficients is 0.76 and net exports are 441 barrels/day per 1,000 people.

Figure 5: Increase in Posterior Probability in Response to Oil Shock, by Country

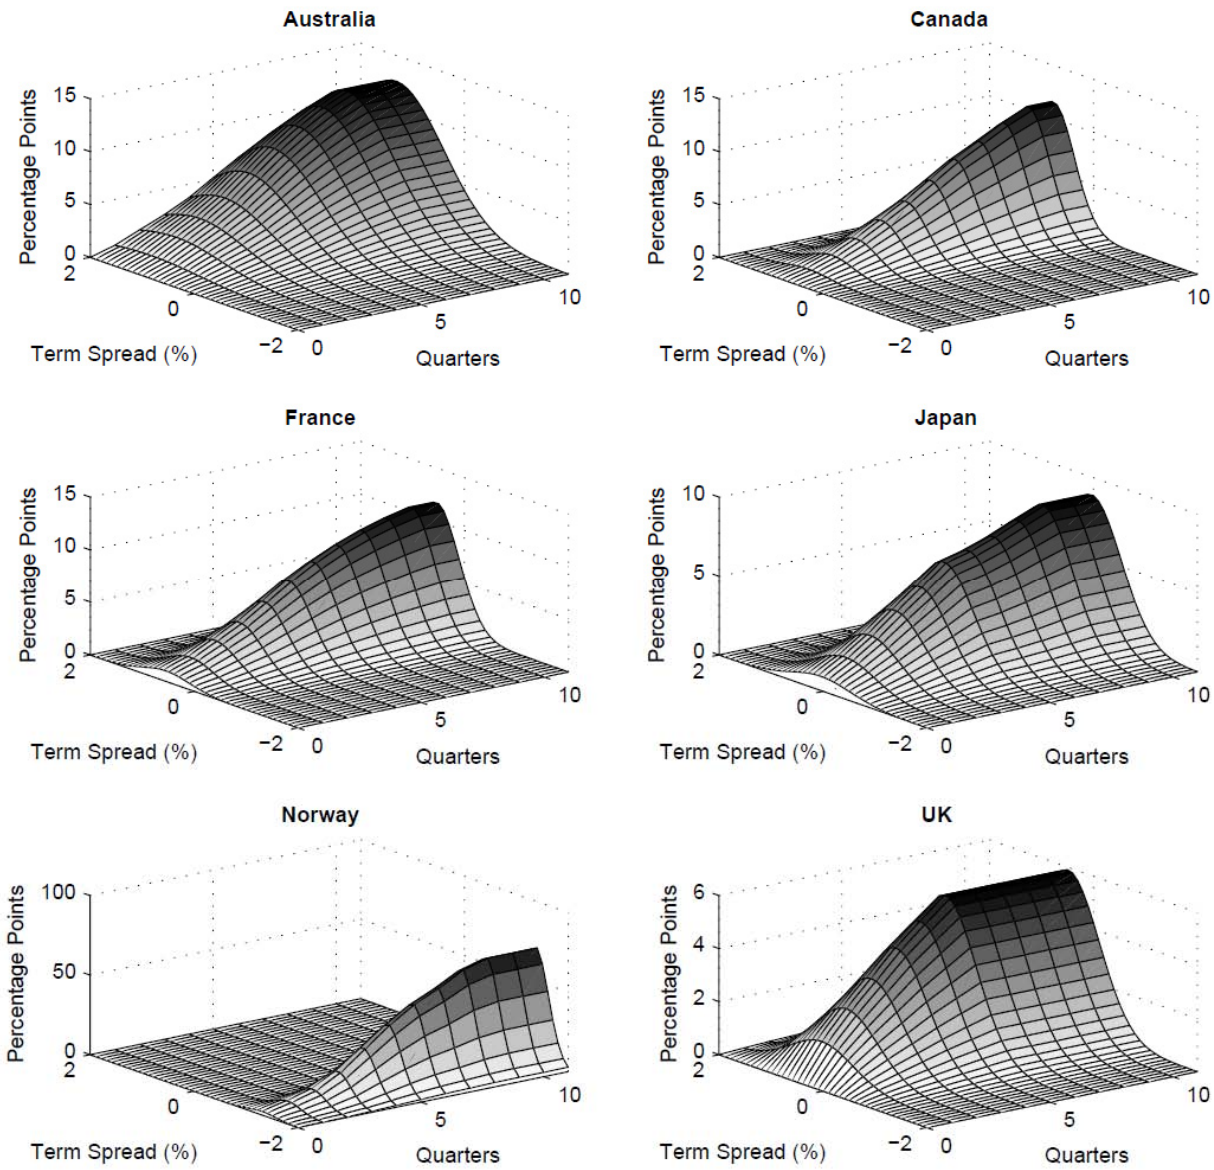

NOTES: These counterfactual responses are calibrated using the preferred specification in Table 5. Ten lags of the world oil series are used for Canada and France; nine for Australia, Japan, and Norway; and six for the UK. Contemporaneous NOPI is included for France, Japan, Norway. The cumulative effects of a one-percentage-point NOPI shock are shown.