

Estimating the Effects of Permanent Oil Price Shocks Consistent with Optimal Factor Allocation

Emmanuel De Veirman and Felipe Labbé *

February 24, 2011

Abstract

This paper estimates the real effects of oil price shocks in New Zealand. In a methodological innovation, we assume that the long-run impulse-responses are consistent with firms' optimal choices about factor allocation. We treat the relative price of oil as an exogenous cointegrating variable, applying the technique of Pesaran, Shin, and Smith (2000). Our results suggest that a single one-standard deviation increase in oil prices implies a cumulative decline in real GDP on the order of 0.20 to 0.30 percent. This suggests that episodes with successive oil price hikes, such as the 2007-08 oil shock, have had substantial effects on aggregate output. Our findings suggest that New Zealand households and firms cannot easily vary the degree of oil usage in response to price changes. Our interpretation is that, if the degree of substitutability were higher, oil price changes would imply less sharp fluctuations in nominal oil expenditure, which would tend to stabilize aggregate output in the face of oil price shocks.

JEL codes: C32, E23, E31, E32, Q43.

Keywords: Cost minimization problem, vector error correction with exogenous I(1) variable, small open economy, macroeconomic effects of oil price shocks.

*Emmanuel De Veirman works in the research section of the Reserve Bank of New Zealand. Felipe Labbé, formerly in the modelling section of the Reserve Bank of New Zealand, is currently at the Banco Central de Chile. We thank Craig Burnside, Alfred Haug, Leni Hunter, Leo Krippner, Kirdan Lees, Dimitri Margaritis, Fabio Milani, Richard Smith, Matthijs van Veelen, Raf Wouters, as well as seminar participants at the Reserve Bank of Australia, the University of New South Wales, Otago University, the University of Auckland, Auckland University of Technology, the New Zealand Ministry of Transport and the Reserve Bank of New Zealand for comments. This paper does not necessarily reflect the views of the Reserve Bank of New Zealand or the Banco Central de Chile. Any errors are ours. Corresponding author: emmanuel.deveirman@rbnz.govt.nz, Reserve Bank of New Zealand, Wellington 6140, New Zealand.

1 Introduction

In this paper, we estimate the effect of oil price changes on New Zealand's real economy.

The international empirical literature on the macroeconomic effects of oil prices typically uses Vector Auto Regression (VAR) analysis to track the response of the domestic economy to an oil price impulse. We contribute to this literature by imposing theoretical structure on the long-run responses. In particular, we require that in the long run, agents purchase the amount of oil that they deem optimal given aggregate income and factor prices. To this end, we specify an empirical equation for oil demand, which we motivate mathematically from the firm's decision problem on how to allocate income across input factors so as to minimize its total cost subject to attaining a particular level of production.

Our approach has the advantage that, by imposing a long-run relation between oil prices and the domestic New Zealand economy, we identify the long-run responses without needing to assume any particular long-run level for the relative price of oil. In doing so, we differ from structural macroeconomic models of the business cycle, which typically have to assume that oil prices converge to a particular long-run level. This is undesirable because there is virtually no theoretical or empirical guidance for choosing any long-run level of oil prices. In fact, we document that the relative price of oil is an integrated variable, implying that it is undesirable to assume mean-reversion in oil prices. In this paper, we specifically address the question of how a permanent change in the price of oil affects the economy.

Much of the empirical literature on the effects of oil price changes focuses on the US economy. Since oil prices are endogenous to US aggregate production and US oil demand, it is particularly challenging to properly identify the part of oil price variation that can be treated as exogenous to the US economy. In this paper, we use data for a small open economy, New Zealand, for which we

can safely claim that domestic conditions have no measurable impact on international oil prices. Because of that, we treat oil prices as an exogenous cointegrating variable, implementing the technique developed by Pesaran, Shin, and Smith (2000).

Our results suggest that a permanent ten percent increase in the international price of crude oil implies a cumulative decline in real GDP on the order of 0.20 to 0.30 percent which fully materializes in the first few quarters after the shock. This effect may appear small, but it suggests that episodes of consecutive increases in oil prices, such as the 2007-08 oil price hike, had a substantial effect on New Zealand real GDP. Our results suggest that an oil price increase, and the associated reduction in aggregate income, are both virtually entirely corrected for by a decline in the demand for oil. We find that the long-run own-price elasticity of oil demand is only about one third of that in the United States, while the income elasticity is slightly larger. In combination, these two channels imply a moderate decline in oil demand after a permanent oil price increase. Oil demand achieves this full, moderate effect by about four quarters after the shock.

Our interpretation is that, if New Zealand households and firms had more opportunities to substitute between oil and other energy inputs, or had more opportunities to substitute between energy-intensive and energy-efficient equipment, any given oil price increase would imply a less pronounced change in nominal oil expenditure, which would tend to dampen the real GDP response.

The literature is divided on whether oil prices have substantial macroeconomic effects. Hamilton (1983) was among the first to document that oil price changes had substantial effects on US real GDP growth. According to his results, large upward spikes in oil prices have tended to anticipate real slowdowns by virtue of a proper causal relationship, not mere coincidence. On the other hand, Bernanke, Gertler, and Watson (1997) found that most of the apparent real effects of oil price changes were in fact attributable to the response of monetary policy to changes in oil prices.

In particular, oil price increases were typically met with monetary policy contractions, where the policy response accounted for much of the subsequent slowdown in the real economy.¹

There is no definitive conclusion in this debate. To our mind, this is so because the size of the economy's response to oil price changes plausibly depends on the environment, and likely changes over time. For instance, Blanchard and Galí (2007) document that the share of overall variation in US output and inflation that is due to oil price shocks has declined over time.² Furthermore, they argue that this is not due to a decline in the size of oil shocks, but rather to a decline in the effect for a given size of shock. They point to decreased real wage rigidity implied by more flexible labor markets, better-anchored inflation expectations implied by increased monetary policy credibility, and a decrease in the oil shares in consumption and production as contributing factors behind the decreased influence of oil shocks on aggregate fluctuations.

To explain time-variation in the effect of oil prices on the economy, Kilian (2009) argues that the effect of a change in oil prices depends on its underlying cause. In particular, oil price increases induced by booms in world activity may have smaller and less immediate negative effects than oil price increases induced by supply disruptions and precautionary demand for oil.

Finally, there is also some literature which assumes that oil price increases may affect aggregate output more than oil price decreases do. For instance, Hamilton (1996) reasons that any unexpected change in oil prices increases uncertainty, which in itself tends to increase the option value of postponing investment projects and therefore tends to dampen aggregate demand. Under that hypothesis, the second-moment effect of oil price changes tends to exacerbate the first-moment

¹Hamilton and Herrera (2004) provide counterarguments to Bernanke, Gertler, and Watson (1997).

²Several other papers find that the effect of oil price shocks on the US economy has declined. For instance, Hooker (1994, 1999) finds that oil prices do not Granger cause US GDP growth after 1980, arguably reflecting structural change in the relation between oil prices and real activity. Similarly, Hooker (2002) finds that since the 1980s, oil price shocks have accounted for only a small fraction of the overall variation in the US price level in comparison with the post-WWII period.

effect of oil price increases, while it tends to dampen the first-moment effect of oil price decreases.

We do not consider time-variation or asymmetric responses in the effects of oil prices. Unlike Kilian (2009), we consider the oil market in one country in isolation, and therefore do not model the determinants of international oil price changes. Instead, domestic oil purchasers take international oil prices as given. On the other hand, an important feature of our model is that we endogenize the quantity of oil purchased. In particular, we derive an equation for long-run oil demand as a function of desired output and factor prices.

The remainder of this paper is structured as follows. In Section 2 we discuss through which mechanisms oil prices could affect macroeconomic outcomes, and give background on the New Zealand oil market. Section 3 motivates the long-run oil demand equation from the representative firm's cost minimization problem. Section 4 discusses the Vector Error Correction Model with oil prices as an exogenous cointegrating variable. Section 5 discusses the results. Section 6 concludes.

2 Background and Data

After briefly discussing mechanisms through which oil prices may affect the economy, we discuss the oil market in New Zealand.

2.1 Supply and Demand Effects

Oil price shocks are typically viewed as aggregate supply shocks, given that oil is a factor into production in various sectors of the economy. An increase in oil prices constitutes an exogenous increase in marginal cost. In response, it tends to prod profit-maximizing firms to increase prices. The resulting increase in the price level implies a decline in real wealth and therefore a decline in

spending on aggregate output, by virtue of a wealth effect.³

In addition, to the extent that the existing factor allocation was optimal, changes in oil prices mean that a different allocation now becomes optimal. The increase in oil prices restricts any oil-purchasing firm's budget set which will tend to mean that it has to reduce output. This is especially the case if it has not yet been able to adjust its factor allocation to the new optimum, but holds true even once it attains the newly optimal factor allocation.

In addition to the above-mentioned effects that operate through oil-purchasing firms, there are likely direct effects on aggregate demand, as long as households directly purchase a non-trivial fraction of oil derivatives such as petrol for their own individual purposes. In that case, an increase in the price of oil affects the price of consumers' basket of goods and services and implies a decline in real wealth, depressing aggregate demand. Simultaneously, there may be a substitution effect at play, in the sense that the increase in oil prices may cause households to purchase more non-fuel goods instead of oil.⁴

2.2 The New Zealand Oil Market

Taken together, oil and natural gas account for about 60% of the energy supply in New Zealand. In 2006, about 80% of oil used in New Zealand was imported. The import share has increased from about 50% in 1997, the year in which domestic New Zealand oil production reached its peak.⁵ New Zealand still produces oil, but its largest oil field⁶ produces oil of a different grade than the

³In the presence of nominal price rigidity, in the short run there will also be a substitution effect from those firms that already increased their prices to those that have not yet adjusted. The substitution effect could imply an effect on aggregate output if there is incomplete information or any real rigidity.

⁴In this paper, we only consider linear effects of oil prices on output. Let us briefly note however that some existing work suggests that the relation may be asymmetric. See Mork (1989) for an empirical paper and Atkeson and Kehoe (1999) for a theoretical model. In these papers, GDP responds more to an increase in the oil price than to an oil price decrease of the same magnitude.

⁵See the relevant International Energy Agency report for New Zealand.

⁶By far the most important source of oil and gas in New Zealand is the Maui field in Taranaki.

oil which is consumed in New Zealand, implying that it primarily produces for export. Currently, over 80% of New Zealand oil purchases are for transport.

The New Zealand oil industry was heavily regulated until the late 1980s. However, this changed drastically with the implementation of the free-market principles embedded in the Petroleum Sector Reform Act from 1988 onwards. Among other measures, the Act removed price controls, ceased government involvement in oil refining, and lifted restrictions on importing refined oil products.

In this paper, we use data for the period 1992Q1-2008Q3. The middle panel of Figure 1 reveals that since 1992, the real volume of purchases of oil products by New Zealand firms and households has gradually declined relative to real GDP. In advanced economies, the oil share tends to decline with further economic development. This can happen as technological improvements make it possible to produce the same output with less oil, or as the economy moves from oil-intense sectors to less oil-intense sectors such as services.

In the top panel, the green line is the nominal US Dollar price of crude oil on the Dubai exchange, while the blue line converts the Dubai oil price to New Zealand dollars and deflates it by the New Zealand GDP deflator with base quarter 1996Q1. From this picture, it is apparent that the amplitude of the oil price fluctuations during the 2000s was much larger than those during the 1990s.

The bottom panel of Figure 1 graphs the share of nominal oil expenditure to nominal GDP, which in principle depends on the relative price of oil as well as on the real oil share. Empirically, we find that fluctuations in the nominal oil share closely track changes in the relative price of oil. This suggests that real oil demand is less than unit-elastic, a finding which we confirm formally in our results section. As oil has tended to become more expensive over the course of our sample, the

nominal oil share has tended to increase too. The nominal oil share hovered around two percent in the first half of the 2000s, but then increased to about 3.5% at the end of the sample. The latter figure is comparable to the oil share in the United States: in 2009, the nominal oil share evaluated at crude oil prices was 4% in the United States.

While the overall size of oil use is similar, our sense is that the nature of oil use is rather different, and it implies that New Zealand households and firms have less leeway to vary the amount of oil they consume in response to price changes. In our results section, we find that the own-price elasticity of oil demand in New Zealand is only one-third as large as that in the United States. We can illustrate the sense that oil is less substitutable in New Zealand by the fact that the average car in New Zealand is twelve years old, which is in line with the observation that households are often not sufficiently wealthy to switch to newer, more fuel-efficient cars. Also, much of the use of heavy utility vehicles occurs in agriculture and forestry, and less so as a leisure activity. Farmers and foresters are plausibly less prone to vary the intensity of use of these vehicles since doing so directly affects their production. As another reason why oil substitution is likely challenging, note that the distances between New Zealand cities are large. These distances imply that road transport uses much fuel. The distances also make it less feasible to build extensive infrastructure for alternative means of transport such as trains.

In this paper, we use two oil price measures. In addition to the Dubai price of crude oil, we also compute a measure of the retail price of oil products in New Zealand. We compute it as a weighted average of the retail prices of regular petrol and diesel, obtained from the New Zealand Ministry of Economic Development. The top panel of Figure 2 graphs the weighted petrol and diesel price. The retail price of oil tends to be less volatile than international crude oil prices. One can see this to some extent in the top panel of Figure 3, which compares relative oil price

inflation for crude oil and retail prices. In our results section, we mention that this translates into the variance of the oil price shock appearing only half as large when measured with retail prices than measured with Dubai prices. This plausibly reflects the fact that in addition to crude oil prices, the retail price includes refining and transport costs, which are plausibly more stable. The retail petrol and oil price also includes GST, which has been unchanged at 12.5% in New Zealand over the course of our sample.⁷

The retail price of oil has increased less steeply over time than crude oil prices have. This becomes most apparent when we turn to the bottom panel of Figure 4, which plots the Dubai and retail prices in the same units. We see that the crude oil price accounted for about one third of the total retail price in the beginning of the sample, but accounted for one half at the end of the sample. This suggests that while the relative price of crude oil price has increased, the margin between retail and crude prices has remained fairly stable in dollar terms. In part, it is also due to the fact that the use of diesel has increased relative to petrol purchases, with diesel being the less expensive of the two fuels. The share of diesel in the weighted sum increases from 35% in 1992 to 50% in 2008.

The share of petrol and diesel consumption in the total demand for oil derivatives is fairly stable at about 80%. One can see this by comparing the middle panel of Figure 1 with the middle panel of Figure 2, or instead by comparing the blue and green lines in the middle right panel of Figure 4.

The bottom panel of Figure 2 reveals that the nominal oil share appears larger when computed with retail prices, which is because retail prices include GST as well as refinery and transport costs.

⁷The GST was introduced in October 1986 at a rate of 10%, increased to 12.5% in July 1989, and increased to 15% in October 2010.

2.3 Other Data Series

We now discuss data sources and graph all data we use for our empirical work. As we will discuss in Section 3, our empirical long-run relationship, equation (6), can be understood as an equation for oil demand as a function of aggregate income, oil prices, and the factor prices of capital and labor. Figure 4 graphs the five variables entering the long-run relation. We deflate GDP, labor earnings, the capital price index, and oil prices by the production-based GDP deflator with base quarter 1996Q1. Oil usage is in millions of barrels. As we discuss in our results section, we find that each of these series is integrated of order one.

We obtain data on the real volume of total purchases of oil derivatives as well as data on the sum of petrol and diesel fuels from the New Zealand Ministry of Economic Development (MED). The middle right panel of Figure 4 graphs those two series. We obtain data on the Dubai price of crude oil from Datastream, and compute the weighted retail price of petrol and diesel from MED data. We will estimate our model with two data sets: one involving the Dubai price for crude oil and total usage of oil derivatives (the blue lines in the middle right and lower panels of Figure 4), and one involving the retail price as well as the total quantity purchased of petrol and diesel combined (the corresponding green lines).

The top left panel of Figure 4 graphs the production-based measure of real GDP in billion 1996Q1 New Zealand Dollar, based on nominal GDP from the Statistics New Zealand National Accounts release. The top right panel graphs the real wage, based on the nominal wage defined as average hourly earnings in ordinary time from the Quarterly Employment Survey. The middle right series is the relative price corresponding to the capital goods price index from Statistics New Zealand. The two most important components of this price index for investment goods are: plant, machinery and equipment; and transport equipment. Our data suggest that investment goods

have become cheaper over time.

Our model for short-run dynamics, equations (10) and (9), involves stationary growth rates in the above five variables. We also enter two stationary control variables. The first of this is relative inflation in commodity prices, shown in the middle panel of Figure 4, computed from the ANZ commodity price index. ANZ is a commercial bank operating in New Zealand. The commodity price index contains prices of New Zealand exports, with important categories being meat, skins and wool; dairy products; horticultural products; forestry products; seafood; and aluminium. We account for export commodity prices because they are typically positively correlated with oil prices. Rises in international commodity prices tend to stimulate the New Zealand economy, as opposed to oil price rises which constitute a negative terms of trade shock. Therefore, it is preferable to single out the effect of changes in oil prices per se, and control for commodity price inflation.

It is also preferable to control for the world business cycle. Oil price fluctuations are influenced by the world business cycle, with an expected positive relation. It is therefore preferable to account for world demand conditions, in order to filter out the direct effect of world demand conditions on New Zealand production, and purely measure the effect of oil price changes per se. To capture world demand conditions, we use the world output gap, computed by RBNZ staff as filtered weighted real GDP of New Zealand's sixteen main trading partners. We plot the world output gap in the bottom panel of Figure 4.

3 Long-Run Relation: Theoretical Motivation

In this section, we motivate a relation involving relative factor prices, oil demand and aggregate output that should hold in the long run as a result of firms' optimal choices regarding factor

allocation.

Consider the representative firm's decision on how to allocate its revenue across three input factors: capital K_t , labor L_t and oil O_t . Assume that the firm operates in perfectly competitive factor markets and has no measurable influence on the aggregate price level, such that it takes the relative price of capital $p_{k,t}$, the real wage $p_{w,t}$ and the relative price of oil $p_{o,t}$ as given. The firm's total cost is then a linear function of the input factors and factor prices: $p_{k,t} K_t + p_{w,t} L_t + p_{o,t} O_t$.

The cost minimization problem formalizes the firm's makes optimal factor allocation decisions. In particular, it shows how the firm chooses capital, labor and oil input in order to minimize its total cost subject to its desire to produce any given level of output \bar{Y}_t . That is to say the firm solves:

$$\min_{\{K_t, L_t, O_t\}} p_{k,t} K_t + p_{w,t} L_t + p_{o,t} O_t$$

subject to:

$$F(K_t, L_t, O_t) > \bar{Y}_t \tag{1}$$

Where $Y_t = F(K_t, L_t, O_t)$ is the production function translating inputs into a single output Y_t .⁸ From the first-order conditions of the cost minimization problem (1), one can derive three equations for optimal factor demand, known as the Hicksian demand for each factor conditional on desired output \bar{Y}_t and prevailing factor prices $p_{k,t}$, $p_{w,t}$ and $p_{o,t}$. We focus on the equation for optimal oil demand. We assume that the form of the production function $F(\cdot)$ is such that the function for optimal oil demand implied by the cost minimization problem is log-linear or can be log-linearized.

⁸Cost minimization is the dual problem to profit maximization. In cost minimization, desired output \bar{Y}_t is treated as given at any given time. We therefore abstract from the firm's decision about how much output to produce. Profit maximization requires cost minimization as well as optimal choice of the level of output.

For instance, when the production function is Cobb-Douglas, the constraint becomes:

$$K_t^\alpha L_t^\beta O_t^\gamma > \bar{Y}_t \quad (2)$$

and the solution for optimal oil demand O_t^* :

$$O_t^* = B \bar{Y}_t^{\frac{1}{\alpha+\beta+\gamma}} p_{0,t}^{\frac{-(\alpha+\beta)}{\alpha+\beta+\gamma}} p_{k,t}^{\frac{\alpha}{\alpha+\beta+\gamma}} p_{w,t}^{\frac{\beta}{\alpha+\beta+\gamma}} \quad (3)$$

Note that with a Cobb-Douglas production function, optimal oil demand is a log-linear function of desired output and factor prices. Taking logarithms, we obtain:

$$\log O_t^* = b + \frac{1}{\alpha + \beta + \gamma} \log \bar{Y}_t - \frac{(\alpha + \beta)}{\alpha + \beta + \gamma} \log p_{o,t} + \frac{\alpha}{\alpha + \beta + \gamma} \log p_{k,t} + \frac{\beta}{\alpha + \beta + \gamma} \log p_{w,t} \quad (4)$$

The cost minimization framework implies not only a negative own-price elasticity for any factor's demand, but also a positive elasticity with respect to prices of the other two factors. That is to say, in the theory of Hicksian demand, inputs are necessarily substitutes. Assuming that each of the Cobb-Douglas coefficients α , β and γ is nonnegative, equation (4) illustrates that this holds in the Cobb-Douglas case. Oil demand relates positively to desired output, which reflects a scale effect: producing more output requires more factor inputs.

Empirically, the assumption that all production factors are substitutes is not necessarily realistic in our case. For instance, intuitively one would expect oil to be a complement to capital goods, out of which the two most important categories in New Zealand are primarily plant, machinery and equipment and transport equipment. In our empirical section, we use a slightly more general

specification for the optimal oil demand equation:

$$\log O_t^* = \beta_1 + \beta_2 \log Y_t + \beta_3 \log p_{o,t} + \beta_4 \log p_{k,t} + \beta_5 \log p_{w,t} \quad (5)$$

That is to say, we do not impose the coefficient restrictions implied by equation (4). More generally, we do not impose that the factors be substitutes.

Equation (5) states that optimized log oil demand is a linear function of log real aggregate output and the prices of the three input factors. In our empirical section, we will assume that equation (4) holds in the long run. Therefore, in the long run, firms purchase the amount of oil that they deem optimal given desired output and factor prices. In the short run however, firms may not always be at their long-run optimal levels of input usage because of the costs involved in adjusting production plans. We express oil demand at any point of time as:

$$\log O_t^d(Y_t, p_{k,t}, p_{w,t}, p_{o,t}) = \beta_1 + \beta_2 \log Y_t + \beta_3 \log p_{k,t} + \beta_4 \log p_{w,t} + \beta_5 \log p_{o,t} + \varepsilon_t \quad (6)$$

where actual oil demand diverges from the long-run relation with production and input factors by a residual ε_t . Theoretically, any deviation from the long-run relation should at some stage be undone, implying that in the long run, actual oil demand equals its optimal value O_t^* given the desired output level and given factor prices. Therefore, in theory the residual ε_t is stationary even if the variables in equation (6) are integrated. Therefore, in principle the variables are cointegrated, a hypothesis that we will test for empirically.

4 Vector Error Correction Model

Now, we describe the Vector Error Correction Model (VECM, or cointegrated VAR) framework which we will use for analyzing how the long-run relation involving optimal oil demand is restored after a permanent change in oil prices. Our approach differs from the standard Johansen (1988) cointegration framework in the sense that we model oil prices as an exogenous cointegrating variable. To do so, we implement the technique developed by Pesaran, Shin, and Smith (2000).

Rewrite the cointegration relation, equation (6), as:

$$\theta' z_t = \varepsilon_t \quad (7)$$

where z_t is a five-by-one vector containing oil demand, aggregate output, and the three factor prices. The vector z_t stacks a vector of exogenous cointegrating variables x_t and a vector y_t of endogenous cointegrating variables:

$$z_t = \begin{pmatrix} x_t \\ y_t \end{pmatrix} \quad (8)$$

In our application, we assume that there is a single exogenous variable -the relative price of oil- entering x_t . The four other cointegrating variables are in principle endogenous and appear in the vector y_t .

The VECM for endogenous variables reads as:

$$\Delta y_t = \gamma + \sum_{i=1}^{p-1} \delta_i \Delta z_{t-i} + \Lambda \Delta x_t + \Pi z_{t-1} + u_{y,t} \quad (9)$$

Which states that every variable's growth rate depends on lagged growth rates of all five variables and contemporaneous relative oil price inflation. Λ is a 4x1 vector while the δ_i 's are 4x5

matrices with short-run coefficients. By virtue of the term $\Pi z_{t-1} = \alpha \theta' z_{t-1} = \alpha \varepsilon_{t-1}$, each of the endogenous variables in principle responds to deviations ε_t from the long-run relation. The vector of adjustment parameters α indicates how quickly, and in which direction, every variable reacts to the error-correction term. $u_{y,t}$ is a vector of reduced-form residuals for the endogenous variables.

While relative oil prices enter the cointegrating relation (6), we assume that they do not respond to deviations from the long-run relation. This captures the fact that oil prices are determined on international markets, and should not measurably respond to economic conditions in a small open economy such as New Zealand. The equation for relative oil price inflation reads:

$$\Delta x_t = c + d \Delta x_{t-1} + e(L) m_t + u_{x,t} \quad (10)$$

Where oil price inflation in principle depends on its own lag as well as a vector m_t capturing stationary variables correlated with oil prices, in particular the world output gap and relative commodity price inflation. Because of the inclusion of this measure of world aggregate demand as well as commodity prices, the residual $u_{x,t}$ reflects variation in oil prices that is orthogonal to the world business cycle and commodity prices.

To estimate the rank of the Π -matrix and to estimate the cointegrating space, we follow Pesaran, Shin, and Smith (2000). Unlike the Johansen (1988) approach which is standard in case all cointegrated variables are endogenous, the Pesaran, Shin, and Smith (2000) approach allows for cointegration between exogenous and endogenous I(1) variables.

In our empirical section, we will assess how shocks to $u_{x,t}$ affect the system. Equation (9) implies that oil price inflation Δx_t affects every of the endogenous variables contemporaneously.

5 Results

Augmented Dickey-Fuller and Phillips-Perron tests for a unit root reveal that all five variables appearing in equation (6) are integrated of order one.⁹ That is to say, we cannot reject the null hypothesis of a unit root in any of the variables, but can reject the hypothesis that the variables are non-stationary in differences.

We select the optimal lag length of the VECM equations (9) by minimizing the multivariate Akaike Information Criterion. Following that criterion, we estimate a VECM with one lag of the growth rates, as well as with contemporaneous oil price growth.

Using the Pesaran, Shin, and Smith (2000) trace and maximum eigenvalue tests for the rank of matrix Π , we detect the existence of a single cointegration relation in the five-variable system.¹⁰ We interpret this relation as the oil demand equation (6).

Table 1 provides the estimated long-run relationships with the Dubai oil price as well as with the retail oil price. In both cases, the long-run elasticity of oil demand to aggregate production is on the order of 0.80. This is consistent with the fact, documented in the middle panels of Figures 1 and 2, that the ratio of real oil usage to real GDP has declined somewhat over the course of our sample. Generally, the output elasticity of oil demand declines with economic development, as firms have access to more energy-efficient equipment and as the sectoral composition moves from oil-intensive sectors to less oil-intensive sectors such as services. In developing countries the output elasticity can be larger than one, while in the United States it is about 0.50.

As expected, the own-price elasticity is negative. It is common to find that oil demand is less than unit-elastic, which implies that oil price increases tend to raise the nominal oil share. By

⁹Table to be added.

¹⁰Table to be added. We test for the cointegration rank by using the critical values in appendix B of Pesaran, Shin, and Smith (2000).

comparing the upper panels of Figures 1 and 2 with their respective lower panel, one readily sees that this implication applies in New Zealand. We find that oil demand tends to decline by 0.02 percent in response to a one percent increase in international oil prices, and by 0.09 percent in response to a one percent increase in retail oil prices. It is common to find smaller elasticities to international oil prices than to retail prices. In both cases, our findings imply that the elasticity of oil demand is only about one third as large as it is in the United States. This suggests that in response to a permanent change in oil prices, New Zealand firms and households have made fairly modest changes to the quantity of oil demanded.

To interpret this finding, note that compared to the United States, much of the oil usage is a necessary part of forestry and agriculture production, or is necessary in terms of transport given the comparatively long distances among New Zealand cities, rather than being a leisure good for which consumption can be easily reduced after an oil price hike. In a similar vein, New Zealand households often do not have the wealth to substitute to newer and more fuel-efficient cars. Finally, a factor behind this result is that we capture the response of oil demand over the course of a twenty-year sample, which may not be sufficiently long to see the full effect of changes in production structure and available technology in response to oil price changes.

Furthermore, note that the elasticity of oil demand to the price of capital goods is -0.15 irrespective of the oil price measure. This suggests that investment goods and oil are complements rather than substitutes. Alternatively, it suggests that the income effect of permanent changes in the price of capital goods dominates the substitution effect. Similarly, the negative coefficient on the real wage suggests that the income effect associated with permanent wage changes dominates the substitution effect.

Figures (5) and (6) document the response of the real economy to one-standard deviation

shocks to the international and retail price of oil, respectively. The international oil price shock is on the order of 10%, which is consistent with the international literature. The retail price shock is on the order of 5%, reflecting the fact that retail prices are less volatile, in line with the fact that besides the crude oil price, retail prices include GST as well as transport and refining costs.

As expected, a permanent increase in oil prices engenders a permanent decrease in real GDP. The cumulative long-run decline is fairly small, on the order of 0.3 percent in the case of the retail price. Statistically, the GDP correction can be small because oil demand responds sufficiently quickly to restore equilibrium. The two main factors which open up an error-correction gap are the initial oil price increase and the resulting decline in aggregate production. The oil price increase tends to call for a decline in oil demand because in this case the substitution and income effects work in the same direction, while the decline in aggregate GDP tends to call for a reduction in factor demand because of a production scale effect. Oil demand does virtually all the error-correction necessary to restore equilibrium.

The real wage declines in response to an oil price shock. Possibly, this happens because of the decline in factor demand associated with the fall in GDP. It suggests that real wages are sufficiently flexible in New Zealand for firms to be able to let the real wage be eroded by inflation in the wake of an oil price increase. We find that the price of capital increases slightly with an oil price shock. This possibly reflects the fact that both oil prices and capital prices reflect international conditions.

Both the response of the real wage and the price of capital have a minor effect on the error-correction term. Our results are robust to omitting the price of capital and real wages from the system.

Taking all these points together, we see that a one-standard deviation shock to oil prices calls

for a moderate decline in oil demand, even accounting for the fact that the oil price increase implies a decline in the scale of production which indirectly calls for a further reduction in oil demand. Oil demand achieves this moderate response sufficiently quickly to restore equilibrium in four quarters. The fact that oil demand makes the adjustment quickly suggests that our estimation primarily captures changes in the intensity of usage of other production factors rather than changes in production plans and technology. The main result is that the GDP response to a single one-standard deviation oil price shock is small: a 10% increase in international crude oil prices implies a cumulative decline in GDP on the order of 0.2%, while a 5% increase in retail prices implies a decline in GDP of 0.3%. This is not to say that oil shocks have little bearing on the real economy. Oil shocks often happen in clusters, and a succession of positive oil price shocks such as those that happened in the 2007-08 oil price hike would have had implied a much larger decline in real GDP.

6 Conclusion

We estimate the effects of oil price shocks on the New Zealand economy using a framework that identifies the long-run responses by assuming that they are in line with firm's optimal decisions about factor allocation. Since New Zealand is a small open economy, we treat the relative price of oil as an exogenous cointegrating variable.

Our results suggest that a permanent one-standard deviation oil price hike implies a cumulative decline in real GDP on the order of 0.2 to 0.3%. This effect is comparable to estimates of the effect of oil prices in the US economy in the last two decades. The effect may appear small, but it does not exclude the possibility that during particular episodes of oil price increases or decreases that last for a few quarters, there is a substantial effect on real GDP. In particular, our finding suggests that the 2007-08 oil price hike has had a substantial effect on real GDP in New Zealand.

In our interpretation, oil demand would adjust in a more pronounced way if firms and households had more opportunities to substitute oil for other energy products or to substitute oil-intensive activities for less oil-intensive alternatives than currently appears to be the case in New Zealand. Such a larger oil demand response would mitigate the effect of oil price changes on the nominal oil share, and therefore reduce the overall effect on New Zealand real GDP.

References

Atkeson, A, and P J Kehoe (1999), “Models of Energy Use: Putty-Putty versus Putty-Clay”, *American Economic Review*, 89, 1028–1043.

Bernanke, B S, M Gertler, and M Watson (1997), “Systematic Monetary Policy and the Effects of Oil Price Shocks”, *Brookings Papers on Economic Activity*, 1, 91-142.

Blanchard, O J, and J Galí (2007), “The Macroeconomic Effects of Oil Shocks: Why are the 2000s so Different from the 1970s?”, *NBER Working Paper* 13368.

Hamilton, J D (1983), “Oil and the Macroeconomy since World War II”, *Journal of Political Economy*, 91, 228–248

Hamilton, J D and A M Herrera (2004), “Oil Shocks and Aggregate Macroeconomic Behavior: The Role of Monetary Policy”, *Journal of Money, Credit, and Banking*, 36, 265-286.

Hooker, M A (1996), ”What Happened to the Oil Price-Macroeconomy Relationship?”, *Journal of Monetary Economics*, 38, 195-213.

Hooker, M A (1999), “Oil and the Macroeconomy Revisited”, Working paper, Board of Governors of the Federal Reserve System.

Hooker, M A (2002), “Are Oil Shocks Inflationary? Asymmetric and Nonlinear Specifications versus Changes in Regime”, *Journal of Money, Credit and Banking*, 34, 540-561.

Johansen, S (1988), “Statistical Analysis of Cointegration Vectors”, *Journal of Economic Dynamics and Control*, 12, 231-254.

Kilian, L (2009), “Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market”, *American Economic Review*, 99, 1053-1069.

Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green (1995), *Microeconomic Theory*, Oxford University Press, New York, NY.

Ministry of Economic Development (New Zealand): Energy Policy Recommendation.

Mork, K A, and R E Hall (1980), “Energy Prices, Inflation, and Recession, 1974–1975”, *Energy Journal*, 1, 31–63.


Mork, K A (1989), “Oil and the Macroeconomy when Prices go Up and Down: an Extension of Hamilton’s Results”, *Journal of Political Economy*, 91, 740–744.

Pesaran, M H, Y Shin, R J Smith (2000), “Structural Analysis of Vector Error Correction Models with Exogenous I(1) Variables”, *Journal of Econometrics*, 97, 293–343.

Rotemberg, J J, and M Woodford (1996), “Imperfect Competition and the Effects of Energy Price Increases”, *Journal of Money, Credit, and Banking*, 28, 549–577.

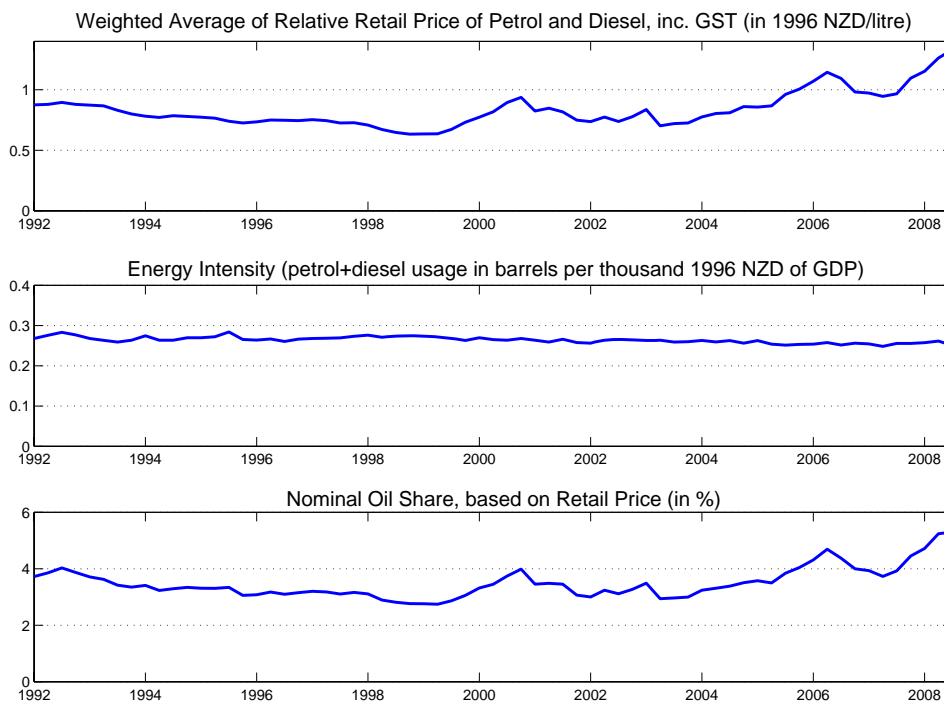
Figures and Tables

Figure 1: International Oil Prices and Oil Consumption in New Zealand



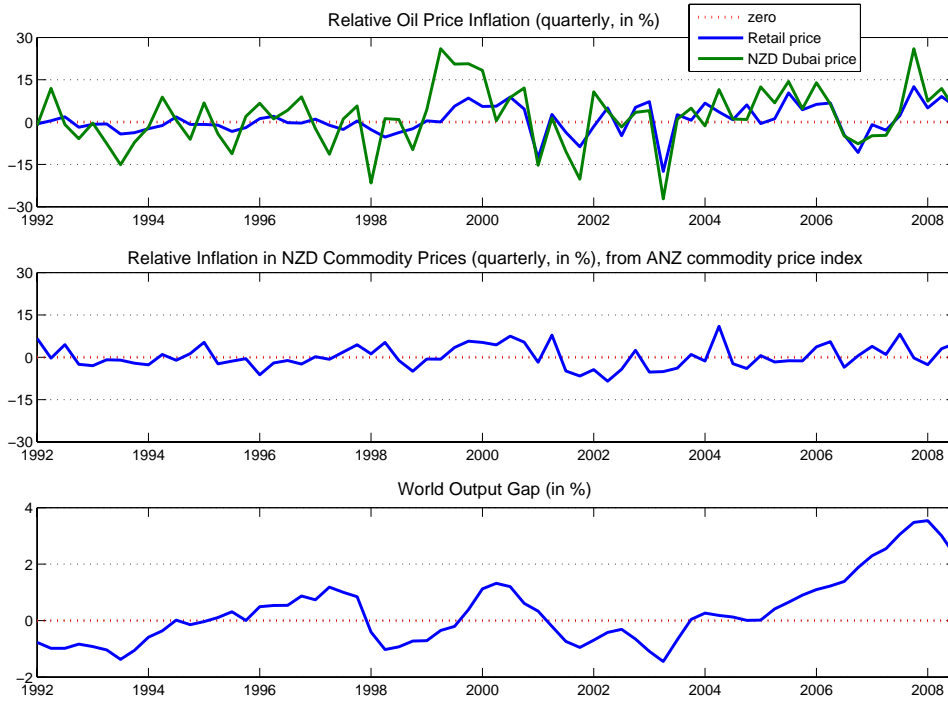
Note: The top panel graphs the relative price of crude oil on the Dubai exchange. The green line plots the nominal Dubai price in US Dollar per barrel. The blue line plots the real Dubai price in 1996Q1 NZ Dollar per barrel. The middle graph plots real oil usage in barrels per thousand 1996Q1 NZ Dollars of GDP. The bottom panel plots the ratio of total nominal expenditure on oil derivatives to New Zealand nominal GDP. The nominal oil share is proportional to the product of the relative crude oil price and real energy intensity. We obtain the Dubai price from Datastream, and data on oil usage from the New Zealand Ministry of Economic Development. For more detail, see Subsection 2.2.

Figure 2: Retail Prices and Petrol and Diesel Consumption in New Zealand




Note: The top panel graphs the relative retail price of fuels in 1996Q1 New Zealand Dollars per litre. We compute the retail price as the weighted average of diesel and regular petrol prices. The retail price includes GST. The middle panel graphs the total usage of petrol and diesel in New Zealand. The bottom panel graphs the implied nominal oil share, computed from quantities of petrol and diesel purchased as well as their retail prices. We obtained retail petrol and diesel prices as well as usage from the Ministry of Economic Development. For more detail, see Subsection 2.2.

Figure 3: Relative Oil Price Inflation and Control Factors



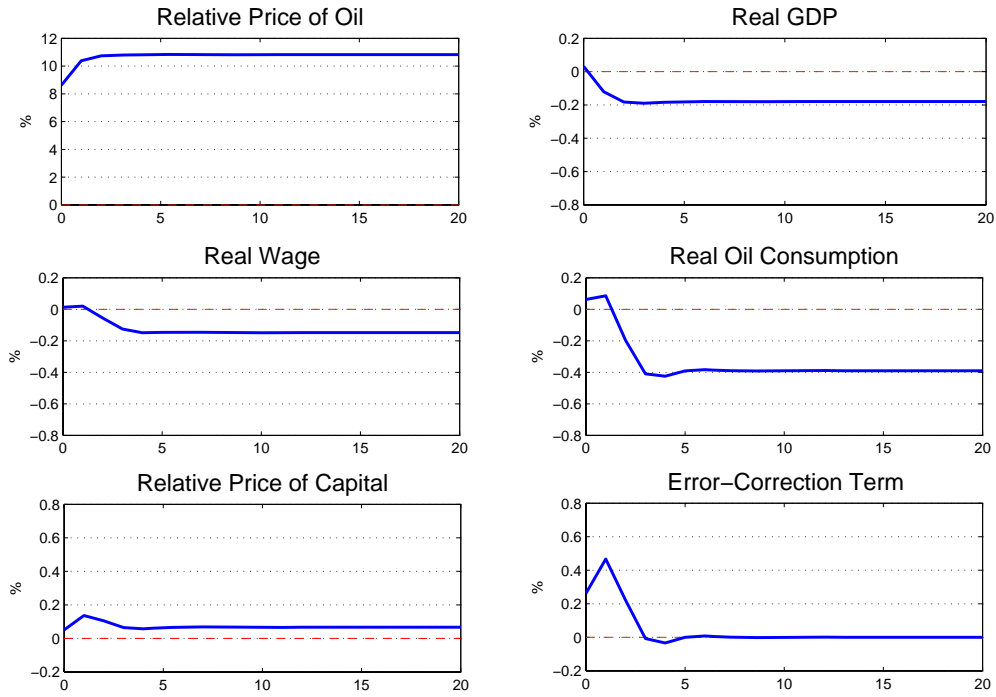
Note: The top panel graphs quarterly relative oil price inflation, in NZD terms, for the retail price as well as for the Dubai oil price, entering equation (10). The middle panel graphs relative inflation in an index of New Zealand’s export commodities produced by ANZ Bank. The bottom panel produces the world output gap computed by RBNZ staff as filtered weighted real GDP of New Zealand’s sixteen main trading partners. For more detail, see Subsection 2.3.

Figure 4: Cointegrating Variables



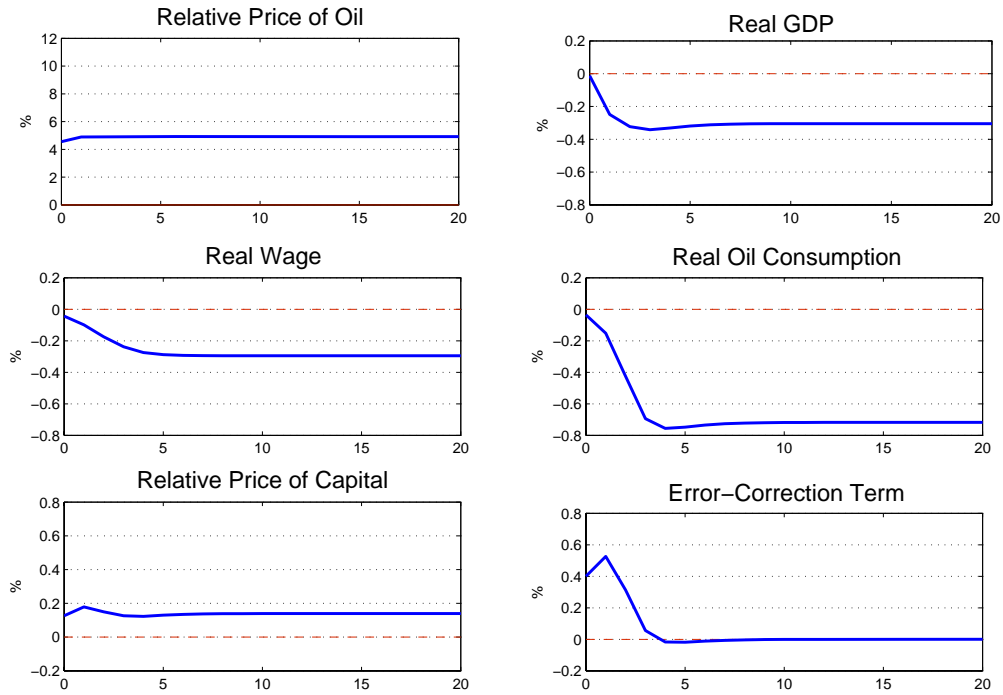
Note: This figure graphs the five variables which enter the long-run oil demand relation, equation (6). All data are in real terms. Oil usage is in millions of barrels. We deflated the four other series by the GDP deflator with base quarter 1996Q1. Gross Domestic Product is from Statistics New Zealand. Average hourly earnings for work during ordinary time is from the Quarterly Employment Survey. The capital goods price index is from Statistics New Zealand. Oil usage and oil price data are from the New Zealand Ministry of Economic Development. We find that all five series are integrated of order one.

Figure 5: Response to Dubai Oil Price



Note: This figure graphs impulse-responses to a one-standard deviation oil price shock to the system made up by equations (9) and (10), with the oil price defined as the Dubai price for crude oil, and oil consumption as the total demand for oil derivatives in New Zealand. The permanent increase in the relative price of oil on the order of 10% implies a cumulative decline in real GDP of 0.20%. Both the oil price impulse and the GDP reaction contribute to opening up the error-correction term, since both call for lower oil demand. Oil demand does virtually all the error-correction. The oil demand response restores equilibrium within four quarters after the shock.

Figure 6: Response to Retail Oil Price



Note: This figure graphs impulse-responses to a one-standard deviation oil price shock to the system made up by equations (9) and (10), with the oil price defined as the retail price for petrol and diesel, and oil consumption as the total demand for petrol and diesel in New Zealand. The permanent increase in the relative price of oil on the order of 5% implies a cumulative long-run decline in real GDP on the order of 0.30%. Both the oil price impulse and the GDP reaction contribute to opening up the error-correction term, since both call for lower oil demand. Oil demand does virtually all the error-correction. The oil demand response restores equilibrium within four quarters after the shock.

Table 1: Cointegration Equations

Retail Price
$\log O_t = -5.75 + 0.87 \log Y_t - 0.09 \log p_{o,t} - 0.15 \log p_{k,t} - 0.03 \log p_{w,t} + \varepsilon_t$
Dubai Price
$\log O_t = -4.60 + 0.81 \log Y_t - 0.02 \log p_{o,t} - 0.15 \log p_{k,t} - 0.24 \log p_{w,t} + \varepsilon_t$

Note: This table shows the estimated long-run elasticities in the oil demand relation, equation (6). We find an income elasticity of oil demand on the order of 0.80. This is in line with a gradual decline in energy intensity displayed in the middle panels of Figures 1 and 2. The estimated own-price elasticities of oil demand are small compared to those in the United States, which suggests a low degree of long-run substitutability of oil in New Zealand. Our findings of a negative elasticity on the price of capital suggests that oil and capital are complements, or that the income effect associated with capital price increases outweighs the substitution effect. Higher real wages also appear to induce lower oil demand.