Oil Prices and the Stock Market*

Robert C. Ready†

First Version: September 1, 2012

This Draft: February 4, 2013

Abstract

This paper develops a novel method for classifying oil price changes as supply or demand driven and documents several new facts about the relation between oil prices and stock returns. Oil supply shocks are significantly negatively correlated with equity returns, and can explain 6% of the monthly variance in the aggregate U.S. Market Return from 1983 to 2012 (10% when the financial crisis is excluded), while demand shocks can explain an additional 38%. The negative effect of supply shocks is not concentrated in industries with heavy oil use, but instead is strongest for consumer goods producers, suggesting that oil shocks act through a restriction on consumer spending. Supply and demand shocks have similar explanatory power for international stock returns, with the strongest effects in oil importing countries. Oil supply shocks are defined as changes in the oil price orthogonal to contemporaneous returns of an index of oil producing firms, with the remaining variation classified as demand shocks. Theoretical and empirical evidence are presented in support of this strategy.

*I would like to acknowledge the helpful comments of Nikolai Roussanov, Ron Kaniel, and Mark Ready. All errors are my own.

†The Simon School of Business - University of Rochester

1 Introduction

Over the past three decades, oil prices have received major attention as an economic indicator.[1] An extensive macroeconomic literature suggests a strong link between oil prices and economic output. Hamilton (2003) and others find a strong negative relation between rises in oil prices and future GDP growth, and Hamilton (2010) points out that 10 of the 11 postwar economic downturns have been immediately preceded by a significant rise in oil prices.

Given these facts, a natural approach is to examine the relations between oil prices and other traded assets, such as equities, to help better understand the links between oil prices and the economy. However in doing this a puzzling fact emerges; oil price changes and stock market returns seem to be unrelated! Indeed, from 1983 to 2012, a simple regression of monthly aggregate U.S. stock returns on contemporaneous changes in oil prices suggests essentially zero relation between the two variables. More simply put: Where is the Oil Price Beta?[2]

This paper attempts to address this puzzle by introducing a novel method of classifying changes in oil prices as demand driven (demand shocks) or supply driven (supply shocks). Supply shocks are defined as the change in the oil price orthogonal to contemporaneous returns of oil producing firms, with the predicted values classified as oil demand shocks.[3] By construction the two series account for all of the variation in oil prices, with supply shocks accounting for roughly 80% of the total variation.

When the two series are examined separately, it becomes clear that the apparent lack of relation between oil prices and stock returns is an artifact of the conflicting effects of the two types of shocks. Instead of no relation, both supply and demand shocks are strongly correlated

[1]As an illustration, a search for occurrences of the term "oil price" in the Wall Street Journal from 1983 to 2012 yields 11,930 articles; more than one article per day.

[2]This disconnect between the observed importance of oil prices for the macroeconomy and the lack of stock market reaction to changes in oil prices has received little attention in the academic literature. The relation between oil prices and stock markets is mostly studied in the context of Vector Autoregressions (VAR), and yields inconsistent results for relations at various leads and lags, and uniformly weak results in terms of contemporaneous correlations. See for instance; Jones and Kaul (1996), Kilian and Park (2009), Chen, Roll, and Ross (1986), Huang, Masulis, and Stoll (1996) and Sadorsky (1999)

[3]In order to avoid look-ahead bias, the shocks are constructed using rolling regressions.

with aggregate stock returns over the sample period, with supply shocks having a highly statistically significant negative relation with stock returns. This relation is also economically significant, with oil supply shocks explaining roughly 6% of the monthly variation in aggregate U.S. stock market returns over the sample period.[4] Additionally, the identified oil demand shocks have a highly significant positive relation with stock returns, explaining an additional 38% of the variation in U.S. stock returns over the same period. Similar results hold for international stock returns, with the impact of supply shocks being strongest in oil importing countries.

The construction of these shocks also allows for an investigation of how oil shocks affect different types of firms, shedding light on the mechanism by which oil shocks affect the economy. The vast majority of industries have strongly negative correlations with oil supply shocks, while the only industries with positive relations are two that intuitively may benefit from high oil prices: gold and coal producers. Surprisingly however, it is not necessarily the industries with the highest oil use that are the most negatively affected by oil supply shocks. While some intuitive industries such as Airlines are near the top of the list, in general producers of consumer goods and services (ie. Apparel and Retail) tend to have greater exposures than high oil use manufacturing firms. This result suggests that the main effect of an oil shock may be a reduction in consumer spending, which in turn impacts firms that produce consumer goods, particularly consumer durables. The least impacted firms appear to be high tech and telecom firms, which have low oil use and may be less directly dependent on consumer spending.

Of course, all of these results hinge on the validity of the identification strategy. To understand the logic, it is important to understand the potential explanation for the lack of correlation between oil prices and stock prices. If an exogenous increase in oil prices is, *ceterus paribus*, bad news for stock prices, what potential explanation is there for the negligible relation between the two variables? One obvious candidate for an omitted variable

[4]The presence of large outliers of returns during the financial crisis increases this to roughly 10% when the sample is restricted to the period prior to June 2008.

is shocks to aggregate demand, which would intuitively be associated with rising oil prices and positive equity returns.[5] In order to account for this and sign the causal relation between oil prices and stock returns, one might potentially find an instrument for oil price changes that is uncorrelated with the rest of the economy, such as a time series of exogenous events affecting oil production.[6] However these events tend to be rare so this technique is less suited to answering the question of how much the variance in stock prices is driven by oil supply shocks.

Another potential issue with this strategy lies in the definition of a supply shock. Although oil supply shocks are typically thought of as discrete events, such as a war or hurricane which disrupts oil production, there is the possibility for more subtle effects. How would one classify a month where oil prices rise slowly day by day, as the amount of oil delivered to ports fails to meet expectations? The resulting change in prices is as much a supply shock as a one time major disruption in production from a natural or political event, but is much more difficult if not impossible to identify from news reports.

Perhaps the most direct technique to account for this problem is examine data on oil production. Kilian (2009) attempts to disentangle demand and supply shocks using VAR with data on oil production and shipping activity as proxies for supply and demand, and Kilian and Park (2009) extend this methodology to examining different shocks' impact on the U.S. stock market. However, they find very little contemporaneous explanatory power (less than 2% combined), mostly concentrated in changes in oil prices related to neither supply nor aggregate demand.

This in some sense is not surprising. Oil production is an incredibly complex endeavor conducted at sites all over the world, and there is no guarantee that the data, when reported, is reported in a timely or accurate manner. Most of this data is not reported at high frequencies, and that which is appears to be fraught with short-term measurement errors which lead

[5]Recent work by Kilian (2009) emphasizes the importance of separating the sources of oil price shocks.

[6]This strategy is pursued by several authors, including Hamilton (1983), Hamilton (2003) and Cavallo and Wu (2006).

to extremely high degrees of mean-reversion. Therefore, while VAR techniques can be useful in observing long run impacts on macroeconomic variables, they may be less appropriate for ascertaining effects on stock markets. This is especially true if some of the variation in oil prices comes from changes in expectations of supply conditions that are never realized. One potential solution to all of these problems is to find a traded asset whose changes in price are correlated with aggregate demand shocks, but unaffected by supply shocks, allowing for its use as an effective control variable. This paper argues that an index of oil producing firms may be just such an asset.

To provide motivation for the identification technique, a theoretical model of a competitive commodity producing sector is introduced, similar to the exhaustible resource models of Carlson, Khokher, and Titman (2007) and Casassus, Collin-Dufresne, and Routledge (2005). The model provides evidence that certain characteristics of commodity production, namely the depletable nature commodity resources, the highly inelastic demand, and the significant difficulties in adjusting levels of reserves, will yield producer stock returns that are unresponsive to changes in the productivity of the sector. This lack of response makes producer returns an effective control for identifying supply shocks.

The basic model does not capture one unique feature of world oil markets; that they include some large nationalized entities that explicitly act as a monopoly. However, when the model is extended to include a monopolistic producer along side a continuum of competitive producers, the returns to the competitive sector of the industry are sufficient to disentangle the sources of demand and supply shocks. This result holds because the shocks to the productivity of the competitive producers tend to generate much longer lasting impact on oil prices than the shocks to the monopolist. This is an interesting result, which suggests a reinterpreting of oil shocks. Rather than headline grabbing events in the OPEC countries, the more important "shocks" to supply may involve the state of production in the large developed economies. Events such as the decline of the North Sea oil fields and the rise of North American oil sands may be much more important to the world economy than unrest in the Middle-East.

Though the model provides a general framework to justify the shock identification methodology, there are other possible explanations. Given that the identification strategy involves regressing changes in the price of oil on returns to oil producers, there is a potential the results are arising mechanically from a source other than oil price shocks. For example, one possible concern is that a change in the market wide discount rate could impact oil producer stock returns without having an effect on oil spot prices. The identification method would then classify an increase in the discount rate as an oil supply shock, which would potentially explain the correlation with the aggregate market return.

To address this concern, the same exercise is conducted with copper and aluminum, two other commodities that are easily storable and require complex production processes, but whose prices are presumably less likely to have a major impact on aggregate output . If the channel for the results is unrelated to oil price shocks, one would expect similar results for these commodities. Instead, while the demand shocks for these two commodities are highly related to the aggregate stock market (and to oil demand shocks), the identified supply shocks have a negligible relation to aggregate stock returns. This is consistent with the two commodities being less fundamental to the economy than oil, but intriguingly the supply shocks do affect certain industries dependent on the good as an input, consistent with a supply shock story.[7]

VAR analysis is also performed to directly test the impact of the identified shocks on variables relating to macroeconomic output and the U.S. oil supply, and the results are consistent with the supply shock interpretation. The identified supply shocks negatively affect current and future economic output and negatively impact current levels of oil consumption and production. Conversely the identified demand shocks positively impact current and future aggregate output, and to the extent they impact levels of oil consumption and production it appears to be in the form of predicting a delayed increase, all results consistent with the

[7]For example, though the industry breakdown is unreported, aluminum supply shocks have the strongest negative effect on producers of soft drinks.

model.[8]

The rest of the paper is organized as follows: Section 2 introduces a basic model of competitive oil producers and discusses the shock identification strategy in the context of the model. Section 3 empirically implements the identification strategy, presents the relations between the shocks and the aggregate stock market, and presents empirical support for the validity of the identification technique. Section 4 presents detailed results on the relations between the two constructed shocks and domestic and international stock markets. Section 5 extends the model to include a monopolist and competitive sector and presents additional empirical results on world oil production of OPEC and non-OPEC countries. Section 6 examines the robustness of the results to different choices of indices and price variables, and illustrates how oil futures returns can be used to construct tradeable strategies to mimic supply and demand shocks. Section 7 concludes.

2 Basic Model

The basic model introduced here is a model of competitive firms which take the price as given and choose both investment in oil reserves (with very high adjustment costs) and the level of a flow input (with low adjustment costs). The model is very standard when compared with previous models of commodity production, such as Kogan, Livdan, and Yaron (2009), Casassus, Collin-Dufresne, and Routledge (2005), Carlson, Khokher, and Titman (2007), and Ghoddusi (2010). The benchmark model views oil as a depletable resource as in Carlson, Khokher, and Titman (2007) and Ghoddusi (2010). The only extension here is that the flow rate is subject to stochastic shocks, in addition to the standard demand shocks, so that the relative impacts of demand and supply shocks on prices and producer stock returns can be examined. Empirically, oil demand is highly inelastic (see Ready (2010)). Given this fact, in this standard setup, it is necessary for oil reserves' rate of depletion to be tied to the flow

[8] Though the model does not explicitly include storage, the two shocks also correlate strongly with inventories, with increases in oil prices from both supply and demand shocks resulting in marked decreases in U.S. inventory levels, which is intuitively consistent with the explanation.

rate in order generate an oil producing industry whose changes in value are unresponsive to oil supply (flow rate) shocks. This unresponsiveness is the property that allows for the identification of the different shocks.

The model does not explicitly describe a claim to an aggregate stock market. While it would be simple to include a reduced form specification of an aggregate market return which would match the observed correlations between aggregate returns on the shocks to the oil industry, this would not provide any additional insight. A more complete model would specify final goods producers which would use oil as an input to produce goods for sale to consumers, but since the main goal here is to provide support for the identification technique, such a model is beyond the scope of the current project.

2.1 Firms

The model consists of competitive firms with standard Cobb-Douglas production technology.

$$O_t = Z_t F_t^{\nu} W_t^{1-\nu} \tag{1}$$

Oil wells W_t, and a flow input F_t, are used to produce oil output O_t. The level of productivity is also affected by an oil industry production shock, Z_t.[9]

In the context of the model, W_t represents oil reserves in the ground. This is an important distinction which is unique to a commodity producer, and helps to capture the storable nature of commodities. The costs of increasing production in a given period are not only the direct costs to the producer of a higher level of the input F_t, but also the reduction of oil reserves available to produce in future periods.[10] This cost is reflected in the evolution of oil reserves

$$W_{t+1} = W_t(1-\delta) - dO_t + I_t \tag{2}$$

[9]For simplicity, it is assumed that there are no firm specific shocks, as well as no entry and exit

[10]This feature is central in the exhaustible resource model of Carlson, Khokher, and Titman (2007), and also present in the production model of Casassus, Collin-Dufresne, and Routledge (2005)

The producer chooses investment in new reserves (I_t) , which are depleted both from depreciation, δ, and from the production of oil multiplied by a depletion rate d. The case where $d = 1$ will be referred to as the "Benchmark Model", and the case where $d = 0$ as the standard "Neoclassical Model".

Given an oil price, P_t, firms sell their output earning a profit Π_t

$$\Pi_t = P_t O_t - c_F F_t - I_t - \Phi(F_t, F_{t-1}, I_t, W_t) \tag{3}$$

Here Φ is a function representing costs to adjusting both the level of the flow input and the levels of investment in oil wells, and has a standard quadratic form

$$\Phi(F_t, F_{t-1}, I_t, K_t) = \frac{a_F}{2}\left(\frac{F_t - F_{t-1}}{F_{t-1}}\right)^2 F_{t-1} + \frac{a_W}{2}\left(\frac{I}{W_t} - \frac{\bar{I}}{\bar{W}}\right)^2 W_t \tag{4}$$

Where $\bar{I}$ and $\bar{W}$ are the deterministic steady-state values of investment and oil well stock, and a_F and a_W govern the level of adjustment costs for the flow input and oil reserves respectively. The producers take the price as given and solve the maximization problem

$$\max_{F_t, I_t} = \sum_{t=0}^{\infty} M_t \Pi_t \tag{5}$$

Where M_t is the stochastic discount factor.

2.2 Consumers

Since the model does not include a separate competitive storage sector (see Routledge, Seppi, and Spatt (2000) and Williams and Wright (1991)), consumers consume all of the oil output in each period. Instead the storable nature of oil is captured in the producers' choice between producing now and saving reserves for future production, and captures some of the same intuition as a separate storage technology.

Consumers of oil are represented by an inverse demand curve, so that spot prices P_t are given by

$$P_t = A_t(O_t)^{-\frac{1}{\alpha}} \tag{6}$$

The price is dependent upon A_t, representing the aggregate level of oil demand in the economy, O_t, the total production of oil, and α, the elasticity of demand.

2.3 Dynamics

From equation (6) there are two possible channels in the model for generating a change in the oil price. The first is a rise in the level of demand A_t, and the second is a reduction in the level of supply O_t. Though producers of final goods and household consumers are omitted for parsimony, simple intuition suggest that rises in the oil price from increases in A_t reflect positive economic news, while rises in price from a reduction in O_t generated by a decrease in productivity, Z_t, would represent negative news for the aggregate stock market.

Both aggregate oil demand and oil productivity are stochastic and their logs (indicated by lower case) evolve according to

$$a_{t+1} = a_0 + \rho_a(a_t - a_0) + \sigma_a e_{a,t+1} \tag{7}$$

$$z_{t+1} = z_0 + \rho_z(z_t - z_0) + \sigma_z e_{z,t+1} \tag{8}$$

Where $e_{a,t+1}$ and $e_{z,t+1}$ are independent normally distributed shocks with mean zero and a variance of one. High realizations of either $e_{a,t+1}$ or $e_{z,t+1}$ correspond to "good" times, and therefore both command positive prices of risk. To capture this the stochastic discount factor is given by

$$\frac{M_{t+1}}{M_t} = \beta \exp(-\lambda_a e_{a,t+1} - \lambda_z e_{z,t+1}) \tag{9}$$

Where λ_a and λ_t are the price of risk associated with the respective shocks, and β is the discount rate.

2.4 Model Results

A competitive equilibrium is defined as a sequence choices of I_t and F_t such that the choice the firms are maximizing firm value while taking P_t as given and the market clearing condition $P_t = A_t O_t^{-\frac{1}{\alpha}}$ is met.

The model is solved numerically, but the main intuition regarding the identification result can be seen in the producers' first order equation for the choice of the flow input, F_t

$$c_F = \nu \frac{O_t}{F_t}(P_t - dq_t) \tag{10}$$

Where q_t is the lagrangian multiplier associated with oil well accumulation, and represents the marginal value of an extra oil well in time $t+1$.

In the standard neoclassical framework ($d = 0$), inelastic demand ($\alpha < 1$) and high adjustment costs on oil reserves W_t mean that a positive productivity shock is detrimental to producers. The increase in output from higher well productivity Z_t yields a decrease in price and a drop in revenue. The competitive firms reduce the level of the adjustable input in response, but the decrease in the input is not enough to counteract the fall in price and therefore results in a reduction in profit.

In the Benchmark Model ($d = 1$), the competitive producers will take into account that selling at the current low price is costly since prices are expected to increase in the future.[11] Therefore the drop in price P_t will come along with a rise in the value of q_t. Accordingly, when faced with a positive productivity shock, the producers will respond with a more aggressive decrease in the level of the adjustable input. The resulting additional decrease in output leads to higher prices in equilibrium, and for reasonable parameters the fall in price will offset the increase in productivity and oil productivity shocks will have no effect on producer stock returns stock returns.

[11]This intuition is a standard feature of exhaustible resource models, dating back to Hotelling (1931).

Table 1: Model Parameters

Parameter	Description	Value
ν	Share of Flow Input in Production	0.6
α	Elasticity of Demand	0.5
σ_a	Volatility of Demand Shock	0.15
σ_z	Volatility of Productivity Shock	0.065
ρ_a	Persistence of Demand Shock	0.95
ρ_z	Persistence of Productivity Shock	0.95
a_0	Mean of Demand Shock	0
z_0	Mean of Productivity Shock	0
c_F	Cost of Flow Input	2.5
a_F	Flow Input Adjustment Cost	3
a_W	Oil Well Adjustment Cost	15
d	Depletion Rate (Benchmark)	1
δ	Deprecation of Oil Wells	0.01
λ_a	Price of Demand Risk	2.0
λ_z	Price of Supply Risk	0.3
β	Discount Rate	0.995

2.5 Simulations

The model is solved for a benchmark case, with parameters given by the Table 1, as well the standard neoclassical case where there is no depletion of reserves from production ($d = 0$). The important parameter for the identification technique is the elasticity of oil demand which is significantly less than one in the data. The calibration uses a value of $\alpha = 0.5$, consistent with other estimates in the literature (See Kogan, Livdan, and Yaron (2009) and Ready (2010)) The remaining parameters are chosen to generally match volatilities and correlations of prices and returns.

To illustrate the mechanisms in the model, Figure figure 1 shows impulse response for the two shocks in each of the two cases. The first column of plots represents the response of four model variables to an oil demand shock. An increase in oil demand generates an increase in oil prices, an increase in oil production, and increased use of the oil flow input by producers.

Table 2: Model Simulated Prices and Returns

Moment		**Data**	**Benchmark Model** ($d = 1$)	**Neoclassical Model** ($d = 0$)
	Observable Moments			
Vol of Oil Price	$\sigma(\Delta p_t)$	0.093	0.092	0.126
Vol of Oil Prod Ret	$\sigma(R_t^{prod})$	0.057	0.062	0.018
Correlation	$\rho(\Delta p_t, R_t^{prod})$	0.612	0.614	0.8808
	Model Productivity Shock Correlations			
Oil Price Changes	$\rho(\Delta p_t, e_{z,t})$	-	-0.691	-0.859
Oil Prod. Returns	$\rho(R_t^{prod}, e_{z,t})$	-	0.014	-0.626
	Model Demand Shock Correlations			
Oil Price Changes	$\rho(\Delta p_t, e_{a,t})$	-	0.612	0.427
Oil Prod. Returns	$\rho(R_t^{prod}, e_{a,t})$	-	0.998	0.759

Monthly volatilities and correlations of innovations to observed and simulated oil prices, (Δp_t), observed and simulated oil producer stock returns (R_t^{prod}), model simulated oil demand shocks $(e_{a,t})$, and model simulated oil productivity shocks, $(e_{z,t})$.

In both cases there is also a positive impact on the value oil producing firms and hence the stock return.

The second column shows the outcome of a decrease in oil well productivity, Z_t. Again, the oil price rises along as production falls, however this time some of the fall in production is offset by an increase in the flow input. This effect is much more pronounced in the benchmark case, as producers compete to produce oil in the time of high prices, before reserves rise and lower prices again. This increase in production is enough to prevent a rise in profitability or value for producers, and hence there is no positive stock return.

Table 2 reports calibrated volatilities and the correlation of oil producers' stock returns and changes in oil prices for both the simulated model and the data. The table also reports the correlations between the simulated price innovations and returns and the unobservable supply and demand shocks. The oil producer stock returns are essentially perfectly correlated with demand shocks while being nearly uncorrelated with supply shocks, making them an effective control for identifying the two unobserved sources of variation.

Figure 1: Model Impulse Response Functions

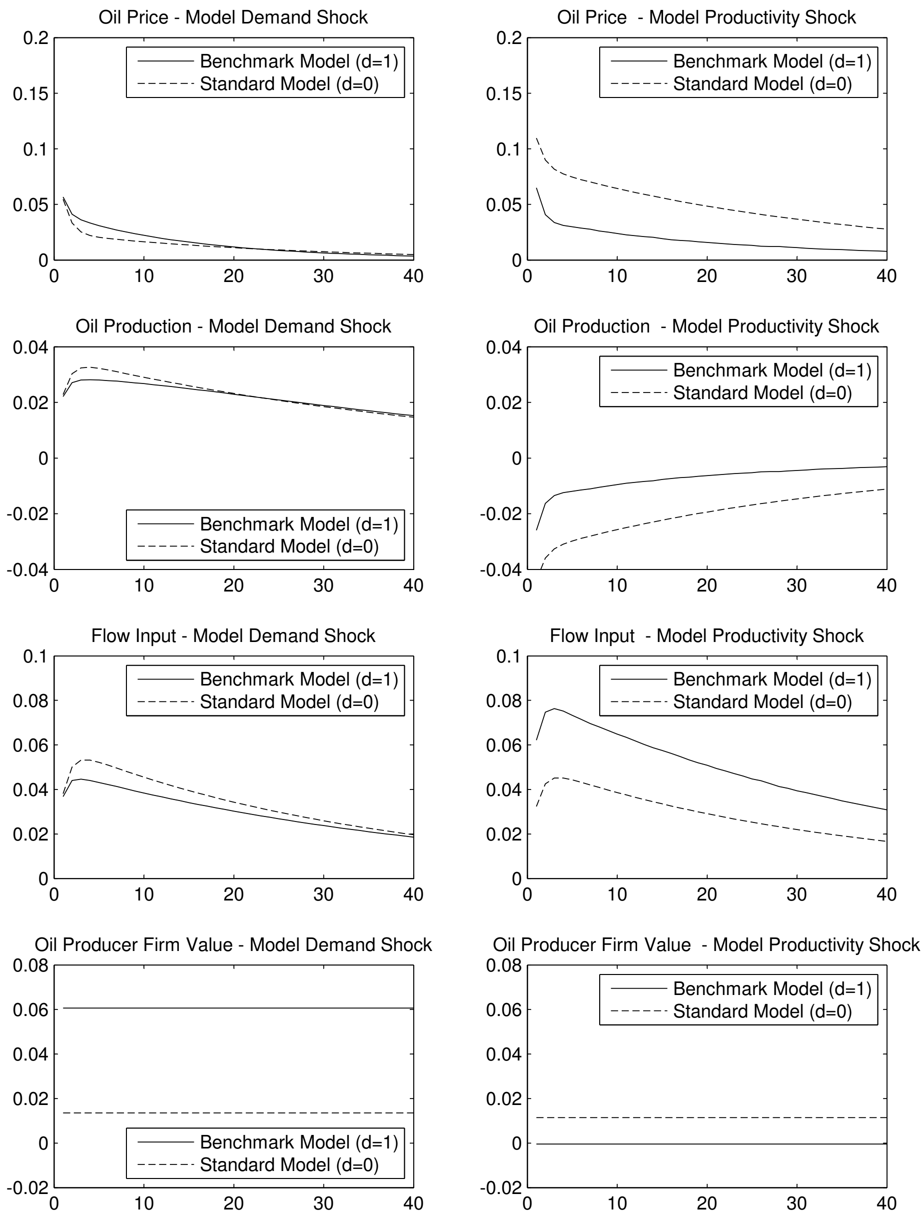

Plots for model impulse responses to demand shocks $e_{a,t}$, and productivity shocks $e_{z,t}$ for two calibrations. In the benchmark case oil production results in a depletion of oil wells ($d = 1$), and in the standard neoclassical case there is no depletion of reserves ($d = 0$).

Though the model is meant to be stylized, it does provide some insight into how an empirical identification technique should work. A positive supply shock should be reflected in lower prices, and higher oil output, while a positive demand shock should be reflected in higher prices, and potentially a rise in oil production, but not one as marked as the observed decrease from the supply shock. The differential exposure of oil producer returns to the two types of shocks also allows for identifying the source of changes in price.

3 Identification of Oil Shocks

The two variables necessary to examine the effect of supply shock are an index of oil producing firms and a measure of oil price changes. To cover as much of the oil production industry as possible, the index used is the World Integrated Oil and Gas Producer Index. This index covers the large publicly traded oil producing firms (Exxon, Chevron, BP, etc..) in most non-OPEC countries, but obviously does not include nationalized oil producers such as Saudi Aramco. The model implications of this are discussed in Section 5. For the change in the oil price, innovations to the log of the nominal value of the West Texas Index (WTI) are used for the bulk of the analysis due to the availability of the longest time series of data. Since the main focus of the paper is on monthly innovations availability of monthly data limits the sample from 1983 to 2012.[12] Aggregate stock market data is the CRSP index. Table 3 provides summary statistics of standard deviations and correlations of the three variables.

This correlation matrix is at the heart of the strategy pursued in this paper. The high correlation between oil producer returns and both oil prices and aggregate market returns, along with a very low correlation between oil prices and aggregate market returns, suggests that there may be some source of variation which loads negatively on aggregate stock returns but positively on oil prices and is uncorrelated with producers, perfectly consistent with the

[12] Monthly data is available from 1979, and 4 year rolling windows are used to identify supply shocks. Section 6 explores the results using other sources of oil producer indices and oil prices, and explores the effects of controlling for aggregate levels of prices and exchange rates. There is so significant impact on results.

Table 3: **Oil Prices Changes, Oil Producer Returns, and the Stock Market (1983-2012)**

	Mean	Stdev	Correlation Matrix		
Aggregate U.S. Stock Returns	0.044	0.162	1		
Oil Producer Index Returns	0.058	0.184	0.604	1	
Oil Price Changes	0.093	0.338	0.059	0.452	1

Annualized means, standard deviations, and monthly correlations. Oil price change change in the log of the WTI. Oil producer returns are from the Datastream World Integrated Oil and Gas Producer Index. U.S. Market return is the aggregate CRSP return. Data are monthly and returns are calculated in logs.

productivity (supply) shocks in the model.

3.1 Constructing Supply and Demand Shocks

In order to study the effects of the portion of oil price innovations which are orthogonal to returns on oil producers, one technique is to simply regress changes in oil prices on contemporaneous returns to oil producers over the whole sample, and then define supply shocks as the residuals from this regression. Although the main results are all qualitatively the same using this simple technique, the approach here will focus on using rolling regressions to calculate the loadings oil price changes on oil producer returns.

This rolling regression is done for several reasons. Rolling regressions avoid any look ahead bias in the construction of the shocks, and also help to control for any changes the loadings that oil producers have on demand shocks, potentially from changes in the level of adjustment costs.[13] The rolling regressions also allow for a single time series of shocks that can then be used in subsample analysis (see Section 4). Finally, when oil futures returns are used as the proxy for the supply shock, this technique allows for construction of tradeable factors which mimic the two shocks (see Section 6).

For the remaining analysis, supply and demand shocks are constructed in the following

[13]Such changes would be closely related to changes in persistence, see Ready (2010) for a detailed discussion of oil price persistence.

manner. Define the demand shock, Δd_t as

$$\Delta d_t = \beta_t R_t^{prod} \tag{11}$$

Where R_t^{prod} is the time t return on the oil producer index, and β_t is the regression coefficient from a regression of changes in oil prices on oil producer returns for the 4 years ending at month $t-1$. The supply shock, Δs_t, is then the residual, which captures the remaining change in the oil price.

$$\Delta s_t = \Delta p_t - \Delta d_t \tag{12}$$

The total annualized volatility for Δp_t over the sample period is roughly 33% while the annualized volatility of the supply and demand shocks is 30% and 15% respectively. Put another way, roughly 80% of the variance in oil prices is classified as supply shocks, with the remaining 20% classified as demand shocks.

3.2 Oil Shocks and Aggregate U.S. Stock Returns

Once the shocks are constructed, it is a simple matter to use them in a basic regression of aggregate stock market returns. Table 4 reports these regressions both for the full sample, 1983 to 2012, and for the sample prior to the financial crisis. The main result of the paper is clear from this table. On their own, oil prices have little or no relation to stock returns. However, when they are decomposed into two series, which are constructed using only time t available information and by construction explain all variation in oil prices, oil prices have a very clear and intuitive link to the stock market. High oil prices from supply shocks are bad news for aggregate stock returns, and can explain 6% of the monthly variation in the aggregate market return, and rises in prices from demand shocks are good news for stocks. The higher coefficient on demand shocks is necessary to account for the lack of relation between oil prices and stock returns, since most of deviation in oil prices is classified as a supply shock.

One potential issue with oil prices, as noted by Hamilton (2003), is that large movements tend to be the drivers of statistical results. To address this Figure 2 shows scatter plots of monthly returns against the realizations of the two shocks against U.S. stock market returns. Although there are some outlying large supply shocks accompanied by negative returns (these tend to be from the gulf war period), the most noticeable outliers are the highly negative returns of the 2008 financial crisis, which tend to go against the general pattern. To illustrate this further, Table 4 also reports results for the supply shock regressions restricting the sample to the period prior to the financial crisis. Doing so provides a marked increase in the R^2 of the supply shocks, from 6% to 10%.

3.3 Evidence for Validity of Identification Technique

It is clear that decomposing changes in oil prices in the manner described in the previous section leads to strong correlations with aggregate stock market returns. In order to defend the interpretation proposed here, that this decomposition allows for identification of supply and demand shocks, several methods are pursued. First as an illustration, the behavior of producer returns and prices around a known supply shock, the first Gulf War, are examined and found to be consistent with the general story. Second, the same exercise is pursued with two other commodities, aluminum and copper, to provide evidence that the result is not an artifact of a purely mechanical relation between producer returns and commodity prices. Finally, VAR analysis is presented which suggests that the two shocks generate patterns in variables relating to macroeconomic output and the oil supply which are consistent with the supply shock interpretation.

3.3.1 Oil Producer Returns and Oil Prices

To defend the validity of this technique, as a first step it is instructive to study the behavior of oil producer returns and oil prices around an observable shock to oil production. Figure 3

Table 4: Aggregate U.S. Stock Market Returns and Oil Price Shocks

Full Sample: 01/1983-12/2012

Dep Var	Constant	Oil Price Change	Oil Demand Shock	Oil Supply Shock	R^2
U.S. Stock Market Returns	0.005**	0.030			0.001
	(0.002)	*0.026*			
	0.002		0.643**		0.376
	(0.002)		*(0.064)*		
	0.005**			-0.133**	0.060
	(0.002)			*(0.039)*	

Pre Crisis Sample: 01/1983-06/2008

Dep Var	Constant	Oil Price Change	Oil Demand Shock	Oil Supply Shock	R^2
U.S. Stock Market Returns	0.006**	-0.030			0.003
	(0.002)	*(0.026)*			
	0.002		0.639**		0.290
	(0.002)		*(0.089)*		
	0.005**			-0.161**	0.105
	(0.002)			*(0.037)*	

Regressions of Aggregate U.S. Stock Market Returns (return to the index of all CRSP stocks) on changes in oil prices and on constructed demand and supply shocks. Demand shocks are constructed as $\beta_t R_t^{prod}$, where R_t^{prod} is the time t return to an index of oil producing firms, and β_t is the coefficient from a prior 48-month rolling regression of changes in oil prices on returns to the index. Supply shocks are then the difference between the time t change in oil prices and the time t demand shock. White (1980) standard errors in parentheses

Figure 2: U.S. Monthly Stock Returns and Oil Shocks

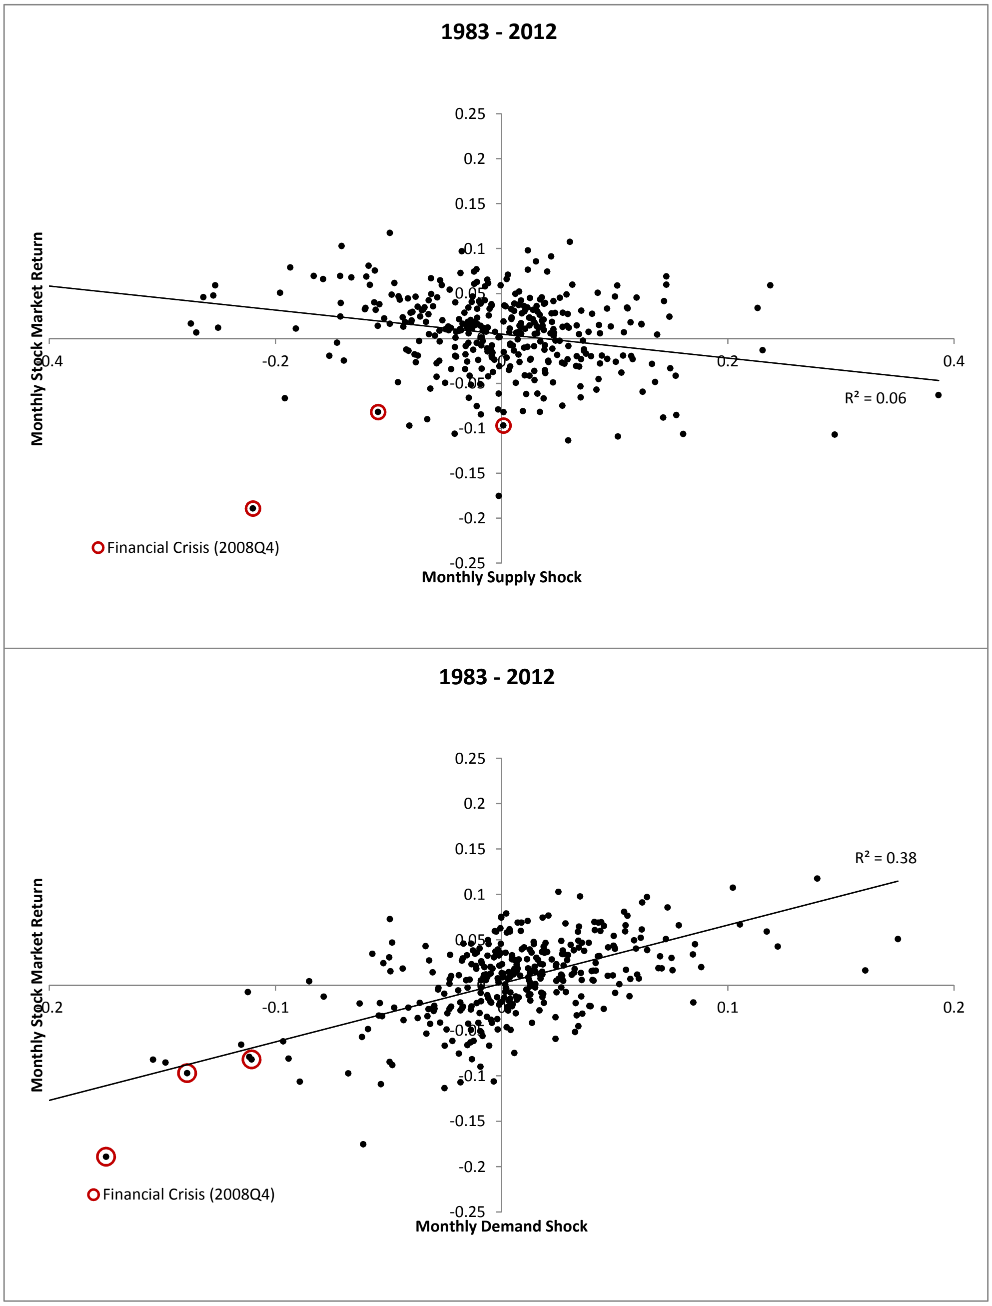

Scatter plots of monthly aggregate stock returns against realizations of oil supply shocks and demand shocks (as described in Table 4). Observations from the three months at the beginning of the Financial Crisis (September - November, 2008) are circled.

gives an example of this using the first gulf war.

In the first panel, the stock prices of several multinational oil producing corporations are shown along with the spot price of oil during the first Gulf War. Companies active in the Middle East, such as Ashland, Inc. and Occidental Petroleum, suffered drops in value due to concerns about their ability to produce during the conflict. In contrast companies, such as British Petroleum, which were not operating in the region, saw their valuations rise with the higher price of crude.

The second panel shows the behavior of the Datastream World Integrated Oil and Gas Producers index over this period, as well as the performance of the U.S. stock market. The initial invasion of Kuwait saw a spike in oil prices, and a negative return for the aggregate stock market, but very little response in the returns of oil producers.

Although some individual producers saw large changes in price, it appears that in aggregate the lower total production and higher price net out, so that oil producers as an industry enjoyed a natural hedge against the potential supply shock. At the same time, aggregate stock values fell, suggesting a potential relation between changes in oil prices and stock returns.

3.4 Copper and Aluminum Producers and Prices

The regression results in Table 4 provide evidence that their may be a strong relation between oil supply shocks and aggregate stock returns. However one potential concern is that the results may be driven by a shock common to aggregate stock returns and oil producer returns, which has no effect on oil prices. The most obvious candidate for this shock is change to the market wide discount rate. However, if this is the case, it is reasonable to think that such an effect would be generated regardless of the commodity being examined.

Tables 5 repeat the regressions for aluminum and copper producers and price changes. Two commodities chosen due to availabilities of returns for indices of producers. Producer data is again worldwide producer indices from Datastream, and price changes are changes

Figure 3: Oil Producers Stock Prices, Oil Prices, and the Gulf War

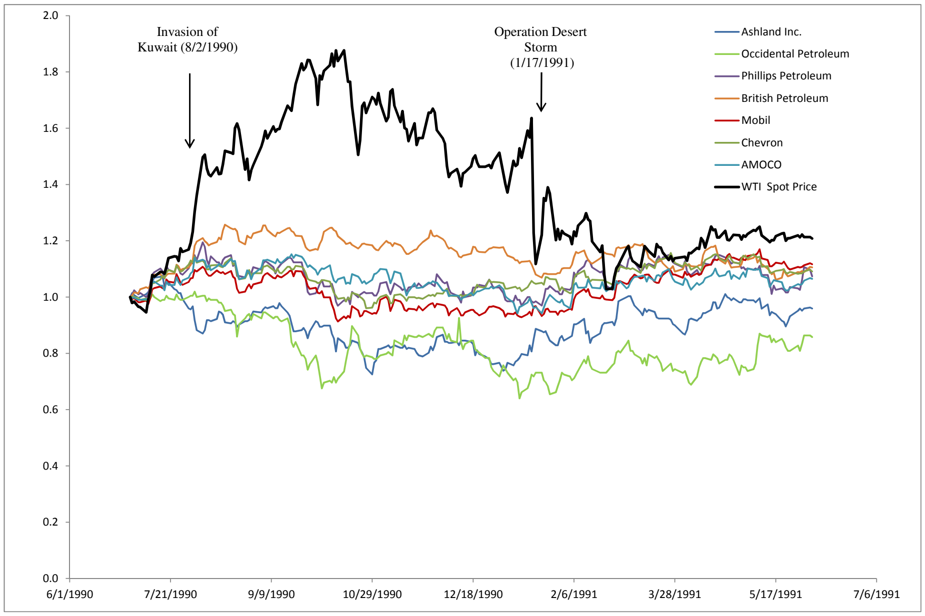

(a) Oil Producer Returns

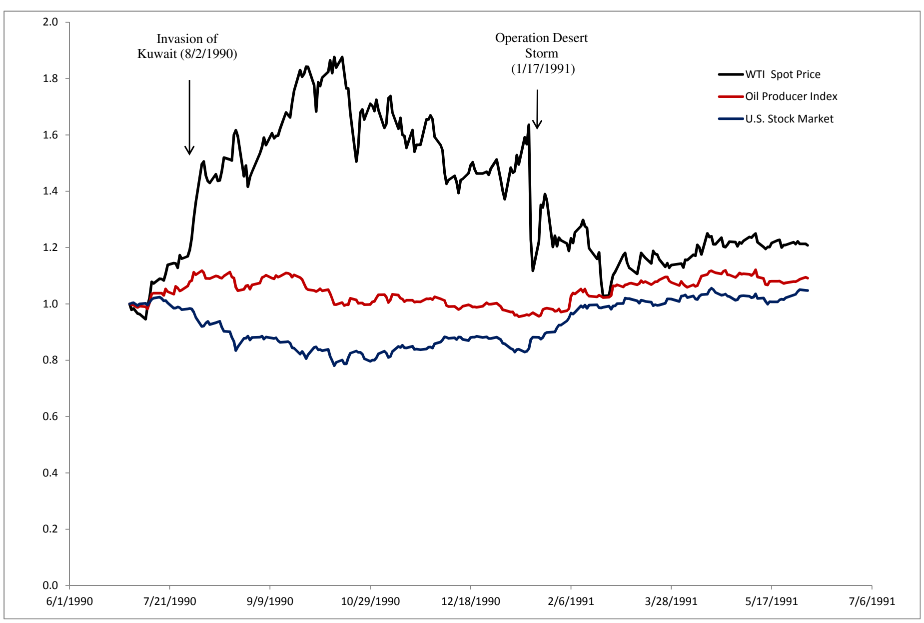

(b) Index Returns

Oil producer index is the Datastream World Integrated Oil & Gas Producer Index. U.S. Stock Return is the value weighted CRSP index.

in the spot prices on the CME metal excahnge. The limited availability of prices leads to shortened samples, so for comparison oil shocks are included over the same sample period. As is clear from the table, the results here are markedly different from those obtained using oil prices. First, both commodity prices have a strong positive relation in a univariate regression with stock market returns. Second, loadings on the supply shocks are not statistically significant, and the R^2 increases are negligible. Since the primary difference between oil and the two metal commodities is the unique status of oil as the most important input commodity for economic output, this result suggests that the driving effect for the patterns in the oil regressions are supply shocks and not some other omitted variable.

3.5 VAR Evidence of Supply Shocks

While the results from regressions of market returns is suggestive of the supply shock story being proposed here, it is important that the identified shocks have an impact on real economic variables, particularly those related to the oil market.

Therefore, to provide a further check of the validity of the identification strategy, Vector Autoregressions (VARs) are estimated for macroeconomic variables, with the constructed supply and demand shocks entering as exogenous variables. Following Hamilton (2008), these VARs have the form.[14].

$$\Delta y_t = \sum_{n=1}^{N} \beta_n^y \Delta y_{t-n} + \sum_{n=0}^{N} \left(\beta_n^s \Delta s_{t-n} + \beta_n^d \Delta d_{t-n} \right) \tag{13}$$

Where y_t is the log of an economic variable of interest, and s_t and d_t are the innovations to the supply and demand indices as described in the previous section. These indices are included contemporaneously as well as with N lags.

[14]Hamilton (2003) estimates a more general form allowing for nonlinearities in the relation, the focus here is on the simpler linear relation, as there appears to be little evidence of any nonlinearity in the relation between oil prices and stock returns. See Figure 2

Table 5: Regressions of U.S. Stock Returns on Copper and Aluminum Shocks

Copper Regressions: 1990 - 2012					
Dep Var	Constant	Copper Price Change	Copper Demand Shock	Copper Supply Shock	R^2
U.S. Stock Returns	0.005**	0.217**			0.132
	(0.002)	*(0.045)*			
	0.002		0.514**		0.347
	(0.002)		*(0.053)*		
	0.005**			-0.044	0.003
	(0.002)			*(0.075)*	

Aluminum Regressions: 1990-2012					
Dep Var	Constant	Alum. Price Change	Alum Demand Shock	Alum Supply Shock	R^2
U.S. Stock Returns	0.006**	0.245**			0.003
	(0.002)	*(0.026)*			
	0.002		1.198**		0.430
	(0.002)		*(0.089)*		
	0.005**			0.010	0.000
	(0.002)			*(0.037)*	

Oil Regressions: 1990-2012					
Dep Var	Constant	Oil Price Change	Oil Demand Shock	Oil Supply Shock	R^2
U.S. Stock Returns	0.006**	0.050			0.011
	(0.002)	*(0.045)*			
	0.002		0.627**		0.360
	(0.002)		*(0.060)*		
	0.005**			-0.105**	0.038
	(0.002)			*(0.040)*	

Supply and demand shocks for the three commodities are constructed in the same manner as Table 4. Producer returns used in construction are returns of the Datastream World Producer Indices for Aluminum and Copper, and Metal Price changes are innovations to the log of the CME spot price. U.S. stock returns are the value weighted CRSP index. Data is Monthly. White (1980) standard errors in parentheses.

Since U.S. data for both aggregate output and oil use is readily available for the sample at issue, I first focus on a set of six variables related to U.S. economic output and oil consumption. The two variables designed to capture the level of aggregate economic activity are real GDP and a Cobb-Douglas aggregate of durable and nondurable household consumption (net of household gasoline consumption). As is shown in Ready (2010), when linking economic variables to levels of the oil price, the closest relation comes from relative levels of this measure of household aggregate consumption and household gasoline consumption. This result motivates the choice of two oil consumption variables, the U.S. total oil consumption as measured by U.S. Oil Production plus Imports minus Exports minus Changes in Inventories all taken from the EIA, and U.S. household gasoline consumption. Inventories the standard measure of U.S. oil inventories (net of the Strategic Petroleum Reserve) provided by the EIA. In addition to aggregate measures of economic activity and utility, it is also interesting to examine the relation between the different shocks and levels of oil inventories. The last variable is the level of international oil production again from the EIA. All data is aggregated to the quarterly level, and regressions are done using 4 lags to be consistent with the convention in the literature.

Figure 4 reports the impulse response functions of six variables to an oil supply shock and Figure 5 reports the same for a demand shock. For an increase in oil prices from a supply shock, there are significant decreases in all six variables, consistent with the hypothesis that this shock measures a reduction in available oil, leading to lower oil consumption, which in turn lowers aggregate economic activity. Conversely, for a demand shock, there are significant positive increases in economic output and total consumption, and an increase in aggregate economic oil use. There is no significant change in household oil consumption, and though it is not significant there is a delayed increase in world oil production, consistent with the model. Also, both supply and demand shocks result in a lower level of U.S. inventories. This is an important result, since a concern with the identification strategy is that the constructed supply shock is simply negatively correlated with the true demand shock. However, if this were the case, all of the VARs should have symmetric impulse response functions, but both with inventories and household oil consumption, this is not the case.

4 Oil Shocks and the Stock Market

The high frequency of the identified supply and demand shocks allow for analysis of the relations between oil prices and the economy that are difficult or impossible using standard VAR techniques. This section explores several of these. The first is to examine the time-varying nature of the relation between oil supply shocks and the economy. The second is to look at the relations between oil shocks and different industries' stock return to better understand how an oil shock effects the market cross-sectionally. Lastly the relation between oil shocks and different countries stock returns are examined. As the subsample analysis shows, the bulk of the correlation between supply and demand shocks occurs pre-crisis. Therefore, when examining industry and country returns, the sample will be truncated at June of 2008.

4.1 U.S. Stock Market and Oil Prices by Sub Period

Several recent papers have called attention to the changes in oil futures market behavior over the period from 2003 to 2008. (see Ready (2010), Baker and Routledge (2011), Hamilton and Wu (2012)) Though the period of available data being used here is short, the monthly frequency allows for enough observations to examine the relation between stock prices and the two oil shocks over smaller sub samples, so in order to see if the relation has changed over the sample period the 1983 - 2012 sample is split into four subperiods. These samples are illustrated in Figure 6. The four samples examined are: 1) 1983 to 1991, covering the oil glut of the mid 1980s as well as the first gulf war, 2) 1992 to 2003, a period of low stable oil prices which also saw the dot-com boom and bust in U.S. markets. 3) 2003 to mid 2008 focused on in recent studies, which saw huge rises in prices as well as worries about

Figure 4: Supply Shock Impulse Response Functions

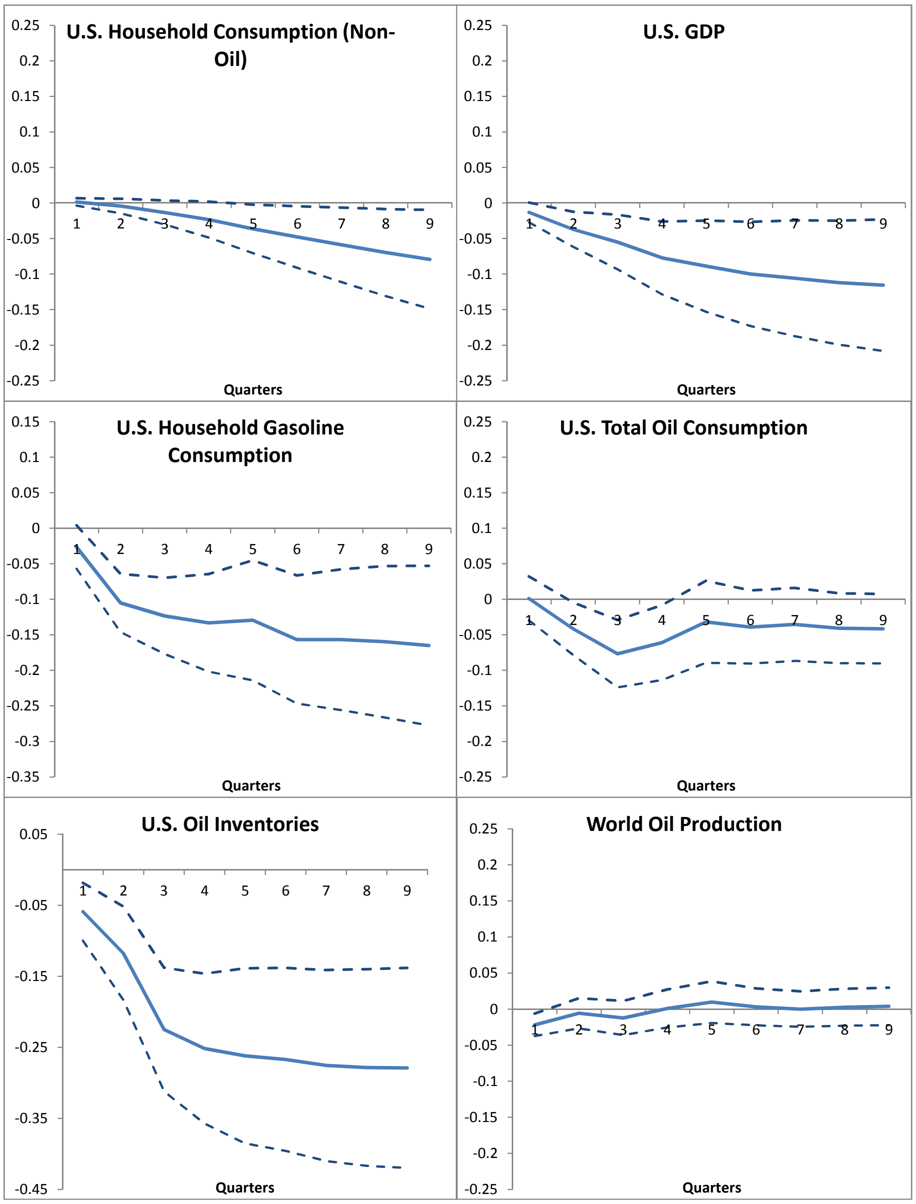

VAR impulse response functions to a one standard deviation oil supply shock for the logs of U.S. real GDP, U.S. total oil use, aggregate household consumption (net of gasoline), household gasoline consumption, U.S. inventories, and world oil production. Regressions performed at quarterly frequency from 1983 to 2012 with four lags.

Figure 5: Demand Shock Impulse Response Functions

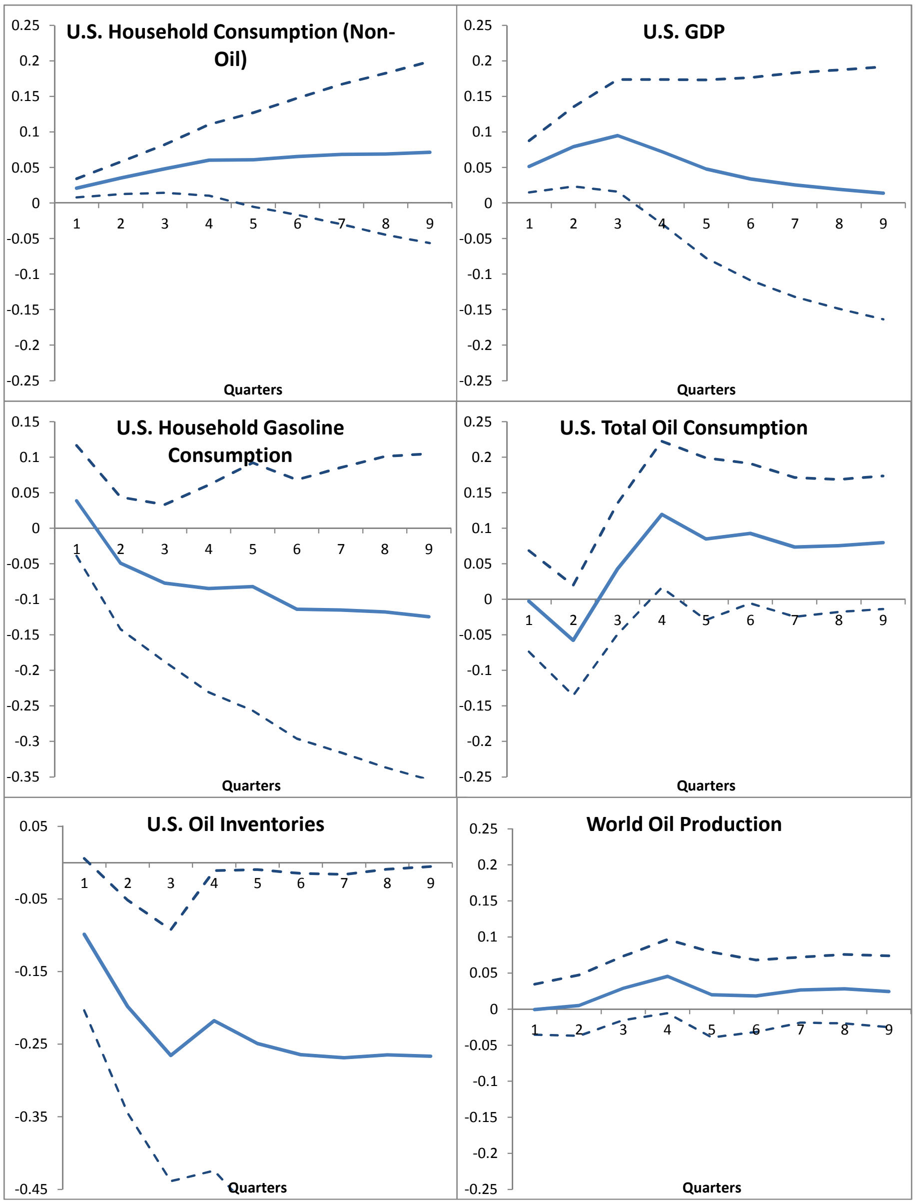

VAR impulse response functions to a one standard deviation oil demand shock for the logs of U.S. real GDP, U.S. total oil use, aggregate household consumption (net of gasoline), household gasoline consumption, U.S. inventories, and world oil production. Regressions performed at quarterly frequency from 1983 to 2012 with four lags.

Figure 6: Oil in the News from 1983 - 2012

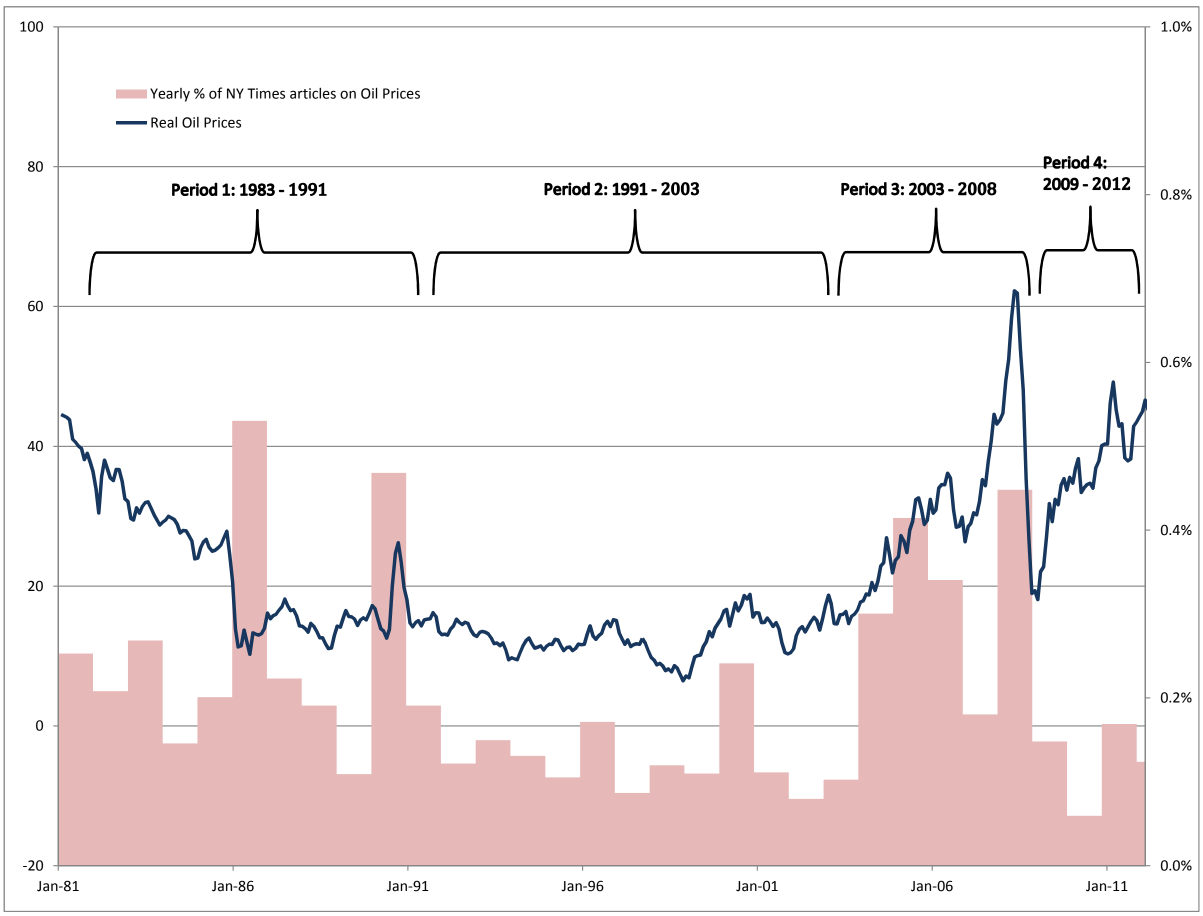

the "financialization" of commodities, and 4) 2009 to 2012, the period post-crisis. Figure 6 also plots the yearly percentage of articles in the New York Times containing the phrase "Oil Prices". As one might expect, the first and third periods were times when oil received the most media coverage, and potentially would be the periods where supply shocks would have the most importance. As the regressions in Table 6 show, this is indeed the case, again providing support for the identification technique, which assigns the highest importance to oil supply shocks during times when the media, and potentially consumers, were most focused on the price of oil.

Figure 7 shows scatter plots for the different subperiods. The first subperiod shows a strong relation between supply shocks and market returns, seemingly driven by large movements in prices, consistent with the large movements around the gulf war and the oil glut of the 1980s. In contrast, the strong relation during the oil price run-up of 2003 to 2008 is driven by a more consistent relation among smaller monthly movements in price.

Another interesting period to note is the post-crisis period from 2009 to 2012. Though this period was accompanied by a precipitous drop in oil prices, they have since risen back up to near pre-crisis levels. Along with this rise has been an uptick in the news coverage relating to oil prices. Finally, this period has also seen an unprecedented diversion in worldwide oil markets. The European Brent crude index, which from its inception in the 1980s through the crisis was nearly perfectly correlated with the U.S. WTI, has since the crisis been trading at premiums of anywhere from 5% to 10% and higher. Additionally the two indices are no longer highly correlated on a monthly basis, with historical correlations of near 0.9 prior to the crisis dropping to below 0.6 post-crisis. Interestingly, though the supply shocks identified using the WTI are negative but insignificant during this period, when the same exercise is done with the Brent index the supply shocks yield a significant relation.[15] This is in line with anecdotal evidence of the WTI being more affected by local transport conditions since 2009, and being replaced by the Brent index as the true barometer of world oil markets.

[15]Results are reported in Section 6

Table 6: Subsample Regressions of U.S. Stock Market on Supply and Demand Shocks

Panel A: Oil Demand Shocks

Dep Var	Period	Constant	Oil Supply Shocks	R^2
U.S. Market Return	**1983-2012**	0.002	0.643**	0.378
		(0.002)	(0.064)	
		Subperiods		
	1983-1991	0.002	0.706**	0.366
		(0.004)	(0.157)	
	1992-2002	0.003	0.740**	0.265
		(0.004)	(0.149)	
	2003 - 06/2008	-0.002	0.493**	0.287
		(0.004)	(0.071)	
	2009-2012	0.011*	0.587**	0.598
		(0.005)	(0.093)	

Panel B: Oil Supply Shocks

Dep Var	Period	Constant	Oil Supply Shocks	R^2
U.S. Market Return	**1983-2012**	0.005*	-0.134**	0.063
		(0.002)	(0.039)	
		Subperiods		
	1983-1991	0.004	-0.217**	0.214
		(0.004)	(0.064)	
	1992-2002	0.006	-0.044	0.007
		(0.004)	(0.045)	
	2003 - 06/2008	0.006	-0.228**	0.220
		(0.004)	(0.058)	
	2009-2012	0.018*	-0.136	0.023
		(0.009)	(0.169)	

Regressions of supply and demand shocks effects on U.S. aggregate market returns over period subsamples. White (1980) standard errors in parentheses.

Figure 7: Oil Shocks and U.S. Stock Returns - Subsample Analysis

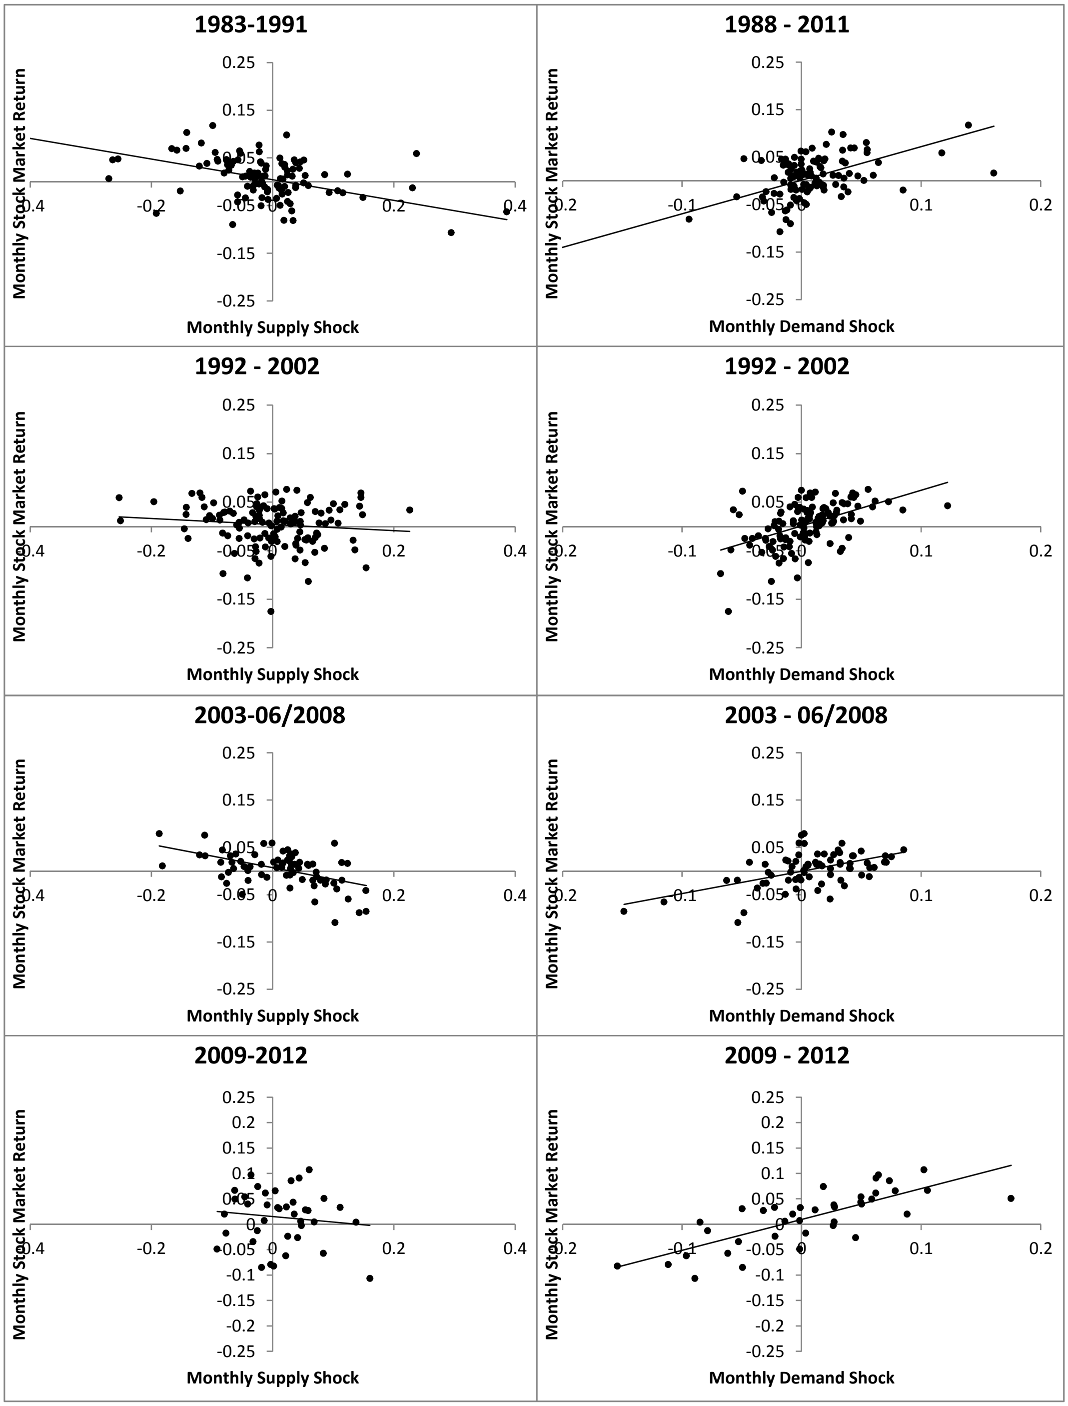

Plots of monthly aggregate U.S. stock returns and constructed oil shocks.

4.2 Industry Portfolios and Oil Shocks

In order to further illustrate the relation between oil prices and stock returns, and to gain more insight into the mechanisms by which oil price shocks impact the economy, Table 7 shows the results of regressing each industry's stock return on both supply and demand shocks. The industries are sorted by their loading on the oil supply shock. What is immediately apparent is that nearly all of the industries do load negatively on the shock. Rather than an oil shock affecting only specific industries that rely on oil as input, nearly all industries suffer when oil prices rise. The notable exceptions being coal companies and gold producers. The positive relation of coal producers is an interesting and intuitive result given its role as a substitute source of energy.

Another interesting feature of this table is that the industries with the most negative shocks are not necessarily the industries one might expect. While some obvious users of oil, such as airlines, are in the top quarter of companies, the top is dominated by financial firms and producers of consumer goods, such as clothing. The companies at the bottom of the list, tend to be manufacturing firms, which also have high oil demand betas and high oil use. To further illustrate this, Figure 8 plots the loading on the supply shock as a function of the relative importance of oil as an input for each industry. This data is calculated using the input output tables from the BEA, and the x-axis of this graph represents the dollars of oil necessary to produce a dollar of output for each industry.

Table 7: Industry Portfolio Returns and Oil Price Shocks (1983 - 06/2008)

Label	**Description**	**Univariate Regressions**						**Oil Use**
		Supply Shock		Demand Shock		Market Return		
		Beta	R^2	Beta	R^2	Beta	R^2	
clths	Apparel	-0.289**	0.160	0.697**	0.163	1.046**	0.555	0.020
		(0.038)		(0.090)		(0.050)		
txtls	Textiles	-0.281**	0.159	0.680**	0.163	0.933**	0.447	0.027
		(0.037)		(0.088)		(0.055)		
airlines	Airlines	-0.268**	0.122	0.569**	0.096	1.064**	0.456	0.068

Continued on Next Page...

Label	Description	Univariate Regressions						Oil Use
		Supply Shock		Demand Shock		Market Return		
		Beta	R^2	Beta	R^2	Beta	R^2	
		(0.041)		(0.100)		(0.062)		
toys	Recreational Products	-0.259**	0.112	0.771**	0.174	1.076**	0.485	0.015
		(0.042)		(0.096)		(0.059)		
rtail	Retail	-0.258**	0.164	0.548**	0.129	1.003**	0.640	0.008
		(0.033)		(0.081)		(0.040)		
autos	Automobiles and Trucks	-0.252**	0.114	0.725**	0.164	1.036**	0.506	0.015
		(0.040)		(0.093)		(0.054)		
insur	Insurance	-0.248**	0.196	0.490**	0.134	0.804**	0.515	0.002
		(0.029)		(0.071)		(0.041)		
banks	Banking	-0.246**	0.147	0.535**	0.122	0.958**	0.573	0.004
		(0.034)		(0.082)		(0.044)		
bldmt	Construction Materials	-0.238**	0.147	0.682**	0.212	1.025**	0.664	0.027
		(0.033)		(0.075)		(0.039)		
persv	Personal Services	-0.235**	0.121	0.672**	0.173	1.026**	0.567	0.008
		(0.036)		(0.084)		(0.048)		
elceq	Electrical Equipment	-0.231**	0.110	0.803**	0.233	1.169**	0.723	0.017
		(0.038)		(0.083)		(0.039)		
beer	Alcoholic Beverages	-0.229**	0.142	0.492**	0.116	0.722**	0.355	0.018
		(0.032)		(0.078)		(0.052)		
fin	Trading	-0.229**	0.100	0.688**	0.159	1.227**	0.767	0.003
		(0.039)		(0.091)		(0.036)		
rubbr	Rubber and Plastic Products	-0.226**	0.126	0.706**	0.215	0.987**	0.606	0.039
		(0.034)		(0.077)		(0.042)		
trans	Transportation	-0.223**	0.134	0.613**	0.178	1.005**	0.622	0.050
		(0.032)		(0.076)		(0.042)		
hlth	Healthcare	-0.222**	0.082	0.725**	0.154	0.945**	0.339	0.008
		(0.042)		(0.097)		(0.070)		
chems	Chemicals	-0.222**	0.143	0.822**	0.342	0.950**	0.651	0.103
		(0.031)		(0.065)		(0.037)		
fun	Entertainment	-0.220**	0.084	0.714**	0.153	1.162**	0.587	0.005
		(0.042)		(0.096)		(0.052)		
paper	Business Supplies	-0.220**	0.142	0.665**	0.228	0.897**	0.571	0.025
		(0.031)		(0.070)		(0.041)		
books	Printing and Publishing	-0.214**	0.141	0.537**	0.155	0.925**	0.646	0.010
		(0.030)		(0.072)		(0.036)		
aero	Aircraft	-0.210**	0.089	0.801**	0.226	1.057**	0.524	0.011
		(0.039)		(0.085)		(0.054)		

Continued on Next Page. . .

Label	Description	Univariate Regressions						Oil Use
		Supply Shock		Demand Shock		Market Return		
		Beta	R^2	Beta	R^2	Beta	R^2	
meals	Restaurants and Hotels	-0.209**	0.123	0.613**	0.186	0.930**	0.556	0.011
		(0.032)		(0.073)		(0.044)		
cnstr	Construction Materials	-0.208**	0.073	0.951**	0.268	1.245**	0.597	0.022
		(0.042)		(0.090)		(0.054)		
boxes	Shipping Containers	-0.203**	0.092	0.748**	0.219	0.885**	0.477	0.032
		(0.036)		(0.081)		(0.049)		
soda	Candy and Soda	-0.201**	0.065	0.566**	0.091	0.716**	0.228	0.014
		(0.044)		(0.103)		(0.070)		
food	Food Products	-0.199**	0.147	0.507**	0.167	0.621**	0.372	0.024
		(0.027)		(0.065)		(0.043)		
fabpr	Fabricated Products	-0.189**	0.062	0.854**	0.220	1.024**	0.437	0.016
		(0.042)		(0.092)		(0.062)		
hshld	Consumer Goods	-0.186**	0.129	0.494**	0.159	0.733**	0.516	0.021
		(0.028)		(0.065)		(0.038)		
labeq	Lab Equipment	-0.183**	0.050	0.845**	0.184	1.374**	0.687	0.010
		(0.046)		(0.102)		(0.049)		
rlest	Real Estate	-0.181**	0.085	0.460**	0.097	0.895**	0.437	0.008
		(0.034)		(0.080)		(0.054)		
softw	Computer Software	-0.179**	0.034	0.759**	0.108	1.648**	0.657	0.005
		(0.055)		(0.125)		(0.063)		
drugs	Pharmaceutical Products	-0.174**	0.095	0.451**	0.112	0.760**	0.479	0.011
		(0.031)		(0.073)		(0.042)		
mach	Machinery	-0.174**	0.062	0.942**	0.320	1.178**	0.724	0.011
		(0.039)		(0.079)		(0.039)		
ships	Shipbuilding / Railroad Eq.	-0.171**	0.049	0.730**	0.157	0.933**	0.352	0.019
		(0.043)		(0.097)		(0.067)		
whlsl	Wholesale	-0.171**	0.092	0.624**	0.216	0.981**	0.700	0.007
		(0.031)		(0.068)		(0.034)		
agric	Agriculture	-0.17**	0.07	0.61**	0.15	0.815**	.366	0.032
		(0.04)		(0.08)		(0.057)		
bussv	Business Services	-0.169**	0.078	0.695**	0.230	1.135**	0.820	0.009
		(0.033)		(0.073)		(0.028)		
guns	Defense	-0.156**	0.048	0.601**	0.124	0.672**	0.209	0.010
		(0.040)		(0.091)		(0.069)		
smoke	Tobacco Products	-0.153**	0.036	0.534**	0.076	0.624**	0.165	0.023
		(0.046)		(0.106)		(0.075)		
chips	Electronic Equipment	-0.153**	0.026	0.755**	0.111	1.503**	0.673	0.009

Continued on Next Page...

Label	Description	Univariate Regressions						Oil Use
		Supply Shock		Demand Shock		Market Return		
		Beta	R^2	Beta	R^2	Beta	R^2	
		(0.054)		(0.123)		(0.056)		
hardw	Computer Hardware	-0.143**	0.023	0.685**	0.094	1.344**	0.572	0.010
		(0.053)		(0.122)		(0.062)		
telcm	Telecommunications	-0.135**	0.054	0.503**	0.131	0.791**	0.514	0.000
		(0.032)		(0.074)		(0.041)		
medeq	Medical Equipment	-0.128**	0.050	0.479**	0.122	0.859**	0.555	0.011
		(0.032)		(0.073)		(0.041)		
util	Utilities	-0.111**	0.063	0.444**	0.177	0.425**	0.234	0.068
		(0.024)		(0.055)		(0.041)		
steel	Steel / Metal Production	-0.111*	0.017	1.066**	0.280	1.297**	0.606	0.023
		(0.048)		(0.098)		(0.056)		
mines	Nonmetallic Mining	-0.094*	0.015	1.054**	0.334	1.007**	0.408	0.052
		(0.043)		(0.085)		(0.065)		
coal	Coal	-0.061	0.003	1.260**	0.227	1.070**	0.241	0.027
		(0.063)		(0.133)		(0.101)		
oil	Petroleum and Nat Gas	0.017	0.001	1.130**	0.708	0.726**	0.351	0.000
		(0.032)		(0.041)		(0.052)		
gold	Precious Metals	0.079	0.004	0.906**	0.101	0.675**	0.076	0.046
		(0.068)		(0.155)		(0.125)		

Table reports regression coefficients and White (1980) standard errors for univariate regressions of 49 Fama-French Industry Return portfolios against oil demand shocks, supply shocks, the aggregate stock market return. The industry "Other" has been replaced by an industry portfolio of U.S. Airlines (SIC Code 4512 in CRSP). The Oil Use column is a measure of dollars of oil input used to create a dollar of industry output constructed using the reported SIC codes from Ken French's online data library and the BEA input output tables.

As mentioned, there does not seem to be a systematic pattern in terms of oil use and oil supply shock beta. Instead, it is companies which are highly dependent on consumer spending, such as clothing, entertainment, and restaurants that seem to feel the pain the most when oil prices rise. This suggests a potential reinterpretation of oil shocks as shocks to consumer spending, rather than a squeeze on oil-thirsty industries.

Figure 9 repeats this plot but this time for oil demand betas. This plot shows that all companies load highly positively on the oil demand shock, and that there does seem to be a relation between oil use and oil demand beta. This suggests that when the manufacturing and oil intensive areas of the economy pick up in activity, the high demand translates to a higher price at the pump for consumers.

Finally Figures 10 and 11 plot the oil supply and demand betas against the aggregate stock market beta of each industry. The graphs clearly show that the oil demand beta is highly correlated with the aggregate market beta, while the supply beta has no systematic pattern. This is further evidence that the demand shock is a fundamental market wide shock.

4.3 International Stock Returns

Up to this point, all of the analysis has been focused on the U.S. stock market. In this section international equity indices are studied to see if the same relations hold true for countries outside the U.S.. Roughly 30 countries are chosen which have equity indices available form Global Financial Data from at least 1988 onward. Table 8 reports oil supply and demand betas for each countries equity index. Notable on this table, are that the world stock index shows nearly an identical pattern of exposure as that of the U.S.. That is a significant negative exposure to oil supply shocks and a highly significant positive exposure to demand shocks.

Table 8: International Stock Returns and Oil Shocks (1987 to 06/2008)

Country	**Univariate Regressions**						**Oil Imports / GDP**
	Supply Shock		Demand Shock		Market Return		
	Beta	R^2	Beta	R^2	Beta	R^2	
Argentina	0.136	0.004	0.940**	0.029	0.752*	0.028	-1.6%
	(0.140)		(0.353)		0.247		
Australia	-0.081	0.011	0.972**	0.290	0.989**	0.358	0.9%
	(0.043)		(0.087)		(0.062)		

Continued on Next Page. . .

Country	Univariate Regressions						Oil Imports / GDP
	Supply Shock		Demand Shock		Market Return		
	Beta	R^2	Beta	R^2	Beta	R^2	
Austria	-0.122**	0.028	0.588**	0.116	0.689**	0.181	1.5%
	(0.045)		(0.100)		(0.090)		
Canada	-0.028	0.002	0.821**	0.270	1.102**	0.596	-1.7%
	(0.050)		(0.097)		(0.065)		
Chile	-0.162*	0.018	0.590**	0.041	0.386*	0.016	4.3%
	(0.069)		(0.164)		(0.151)		
Czech Republic	-0.094	0.008	0.567**	0.052	0.718**	0.098	2.9%
	(0.082)		(0.184)		(0.165)		
Denmark	-0.076*	0.017	0.583**	0.174	0.677**	0.308	-0.8%
	(0.033)		(0.073)		(0.047)		
Europe	-0.156**	0.089	0.718**	0.328	0.976**	0.709	1.6%
	(0.029)		(0.059)		(0.029)		
Finland	-0.096*	0.013	0.617**	0.092	0.752**	0.221	2.0%
	(0.048)		(0.111)		(0.066)		
France	-0.213**	0.100	0.757**	0.238	1.064**	0.580	1.6%
	(0.041)		(0.086)		(0.058)		
Germany	-0.176**	0.059	0.746**	0.186	0.991**	0.442	1.5%
	(0.040)		(0.089)		(0.052)		
Greece	-0.168**	0.025	0.666**	0.068	0.563**	0.074	3.3%
	(0.060)		(0.141)		(0.094)		
Hong Kong	-0.168**	0.030	1.017**	0.195	1.245**	0.246	2.9%
	(0.054)		(0.118)		(0.102)		
India	0.102	0.011	0.317*	0.019	0.229**	0.016	3.9%
	(0.054)		(0.130)		(0.083)		
Indonesia	-0.076	0.003	0.563**	0.033	0.649**	0.053	-0.1%
	(0.075)		(0.177)		(0.159)		
Ireland	-0.236**	0.127	0.653**	0.170	0.744**	0.259	2.0%
	(0.035)		(0.083)		(0.059)		

Continued on Next Page...

Country	Univariate Regressions						Oil Imports / GDP
	Supply Shock		Demand Shock		Market Return		
	Beta	R^2	Beta	R^2	Beta	R^2	
Italy	-0.135**	0.032	0.510**	0.080	0.864**	0.256	1.6%
	(0.042)		(0.099)		(0.069)		
Japan	-0.106*	0.021	0.554**	0.098	0.965**	0.418	1.7%
	(0.042)		(0.096)		(0.053)		
Malaysia	-0.098	0.010	0.862**	0.133	0.981**	0.211	-2.7%
	(0.056)		(0.126)		(0.094)		
Mexico	-0.205*	0.021	1.254**	0.136	0.837**	0.074	-3.2%
	(0.081)		(0.181)		(0.138)		
Norway	-0.022	0.001	1.101**	0.362	1.003**	0.317	-17.9%
	(0.044)		(0.084)		(0.069)		
Netherlands	-0.168**	0.068	0.766**	0.247	1.038**	0.550	2.7%
	(0.036)		(0.077)		(0.054)		
New Zealand	-0.103*	0.017	0.716**	0.144	0.691**	0.197	1.9%
	(0.045)		(0.100)		(0.065)		
Peru	0.063	0.001	0.458	0.012	0.322*	0.009	0.9%
	(0.100)		(0.238)		(0.155)		
Philippines	-0.270**	0.053	0.623**	0.049	0.643**	0.071	4.6%
	(0.065)		(0.156)		(0.109)		
Portugal	-0.159**	0.029	0.618**	0.077	0.787**	0.109	3.0%
	(0.052)		(0.122)		(0.109)		
Spain	-0.223**	0.088	0.674**	0.165	1.133**	0.533	2.5%
	(0.045)		(0.095)		(0.066)		
Sweden	-0.178**	0.058	0.648**	0.136	0.942**	0.395	1.7%
	(0.041)		(0.093)		(0.054)		
Switzerland	-0.205**	0.131	0.569**	0.154	0.839**	0.483	1.2%
	(0.034)		(0.086)		(0.056)		
Taiwan	-0.198**	0.024	0.703**	0.054	0.768**	0.092	6.7%
	(0.072)		(0.169)		(0.113)		

Continued on Next Page...

Country	Univariate Regressions						Oil Imports / GDP
	Supply Shock		Demand Shock		Market Return		
	Beta	R^2	Beta	R^2	Beta	R^2	
Thailand	-0.194**	0.031	0.953**	0.132	0.811**	0.130	5.1%
	(0.062)		(0.140)		(0.106)		
Turkey	-0.113	0.004	1.066**	0.058	1.247**	0.090	2.7%
	(0.116)		(0.264)		(0.243)		
United Kingdom	-0.148**	0.072	0.786**	0.354	1.074**	0.494	1.5%
	(0.030)		(0.061)		(0.051)		
United States	-0.167**	0.123	0.647**	0.324	0.917**	0.753	1.6%
	(0.026)		(0.054)		(0.025)		
World	-0.155**	0.113	0.656**	0.351	1.000	1.000	0.0%
	(0.025)		(0.051)		(0.000)		

Table reports regression coefficients and White (1980) standard errors for univariate regressions of Country Return portfolios against oil demand shocks, supply shocks, the aggregate stock market return. Country returns are from Global Financial Data. Data is monthly from 1987 to 2008.

This relation is not uniform among all countries however, nor would it be reasonable to think that it should be. Clearly, countries which import a large amount of oil, will have a different exposure to these shocks than a country which is not heavily reliant on crude oil imports. To illustrate this, Figure 12 and plots each country's exposure to supply and demand shocks as a function of a measure of each country's dependence on foreign oil. This measure is constructed as the average of net oil imports to GDP over the sample period.

As Figure 12 shows, there is a clear relation between countries which import a large amount of oil relative to the size of their economy and their exposure to oil price shocks. Oil exporters tend to have a strong positive relation to oil supply shocks, while highly advanced economies which are not dependent on oil (ie. Japan) have relatively less exposure. The most exposure is in countries with high oil consumption (and high dependence on cars for transportation), providing further evidence that the channel for oil supply shocks is through

Figure 8: Industry Oil Supply Shock Betas and Oil Use

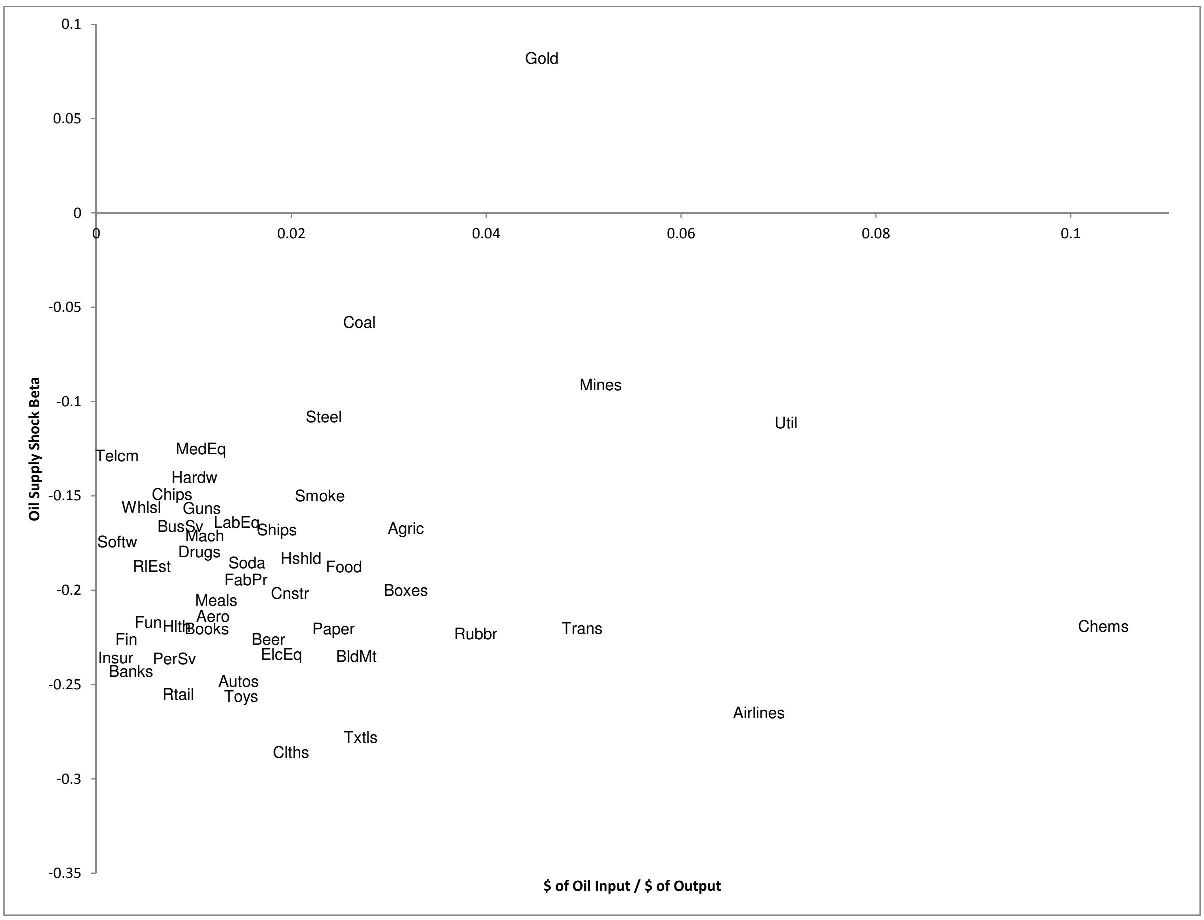

Plots of industry oil use and stock market beta against oil supply shock beta. The X-axis represents the amount of oil in dollars necessary to produce a dollar of output. Data from the BEA input-output tables.

Figure 9: Industry Oil Demand Shock Betas and Oil Use

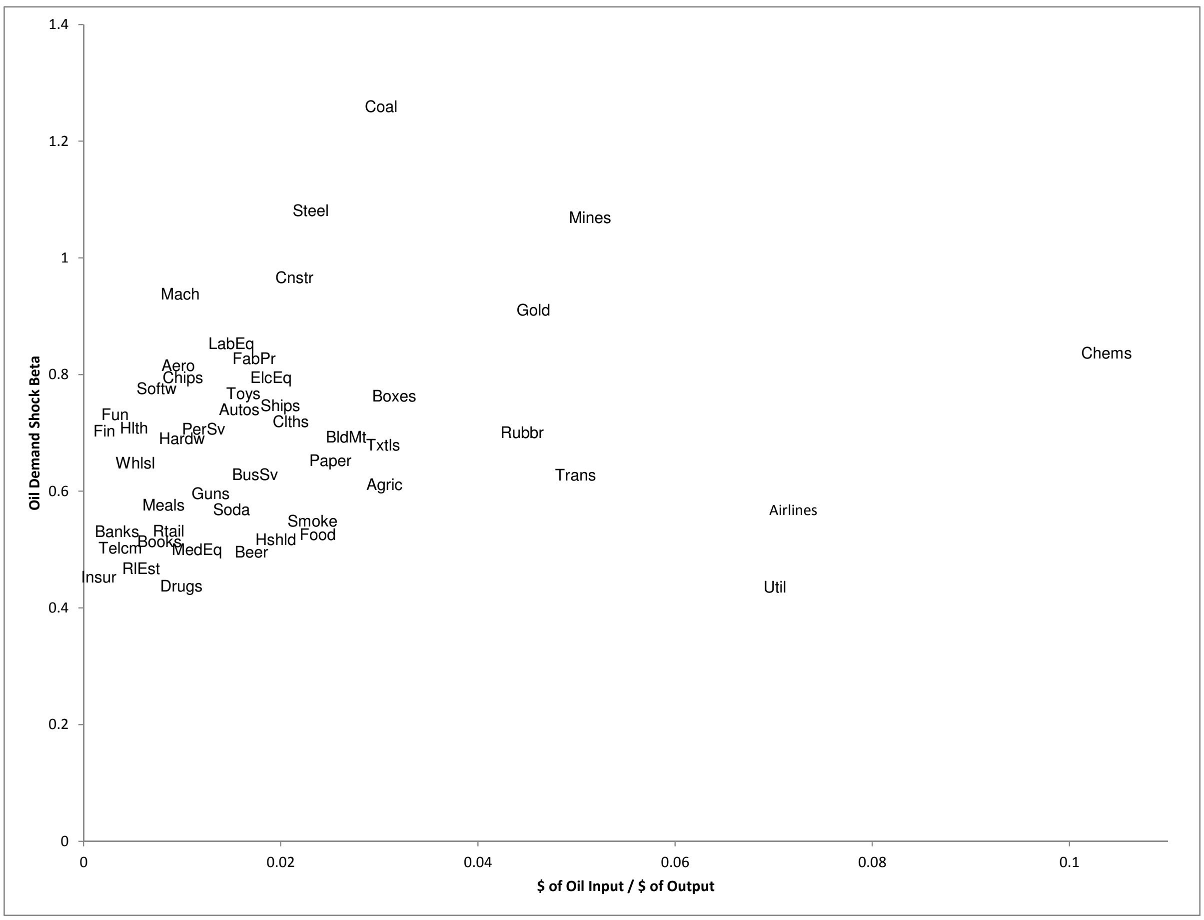

Plots of industry oil use and stock market beta against oil demand shock beta. The X-axis represents the amount of oil in dollars necessary to produce a dollar of output. Data from the BEA input-output tables.

Figure 10: Industry Oil Supply Shock Betas and Market Betas

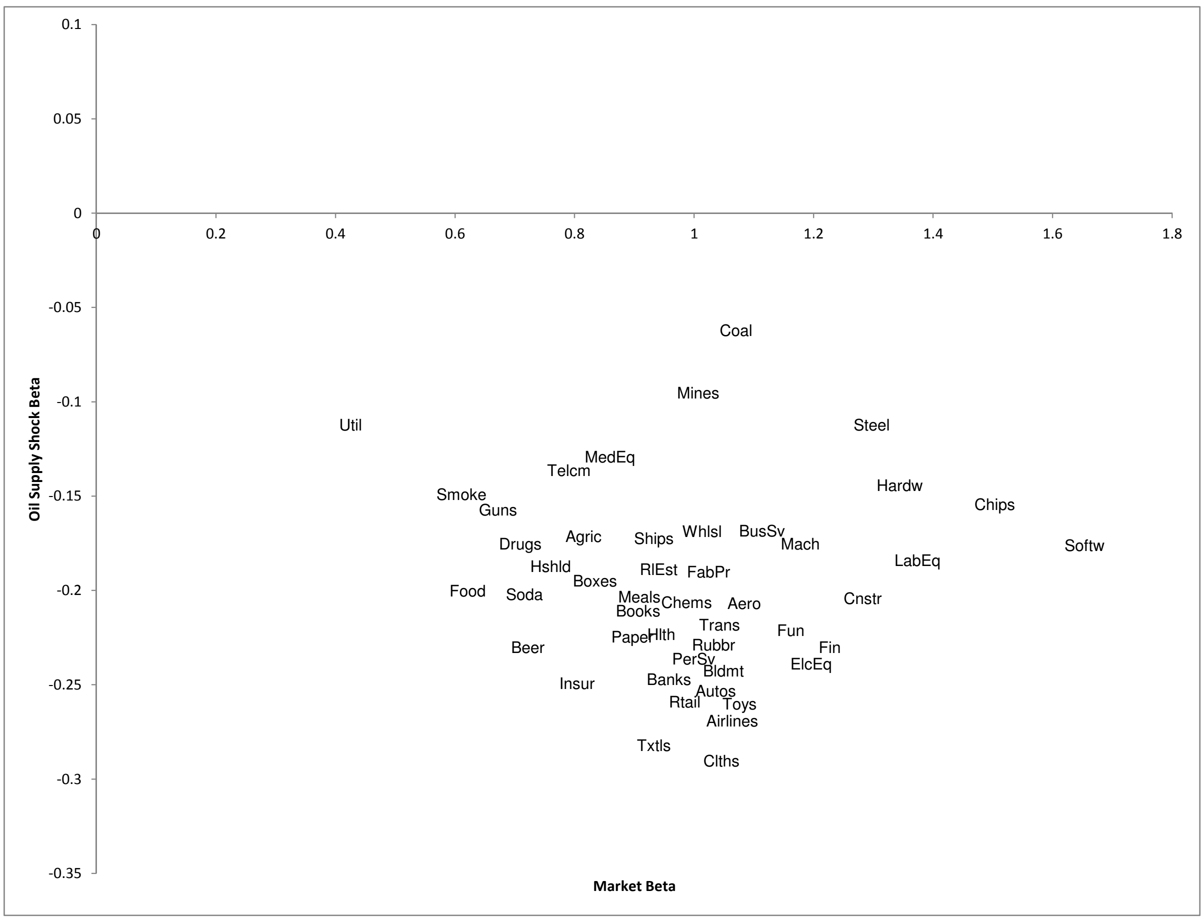

Plots of industry oil use and stock market beta against oil supply shock beta. the X-axis represents the industry beta with aggregate stock market.

Figure 11: Industry Oil Demand Shock Betas and Market Betas

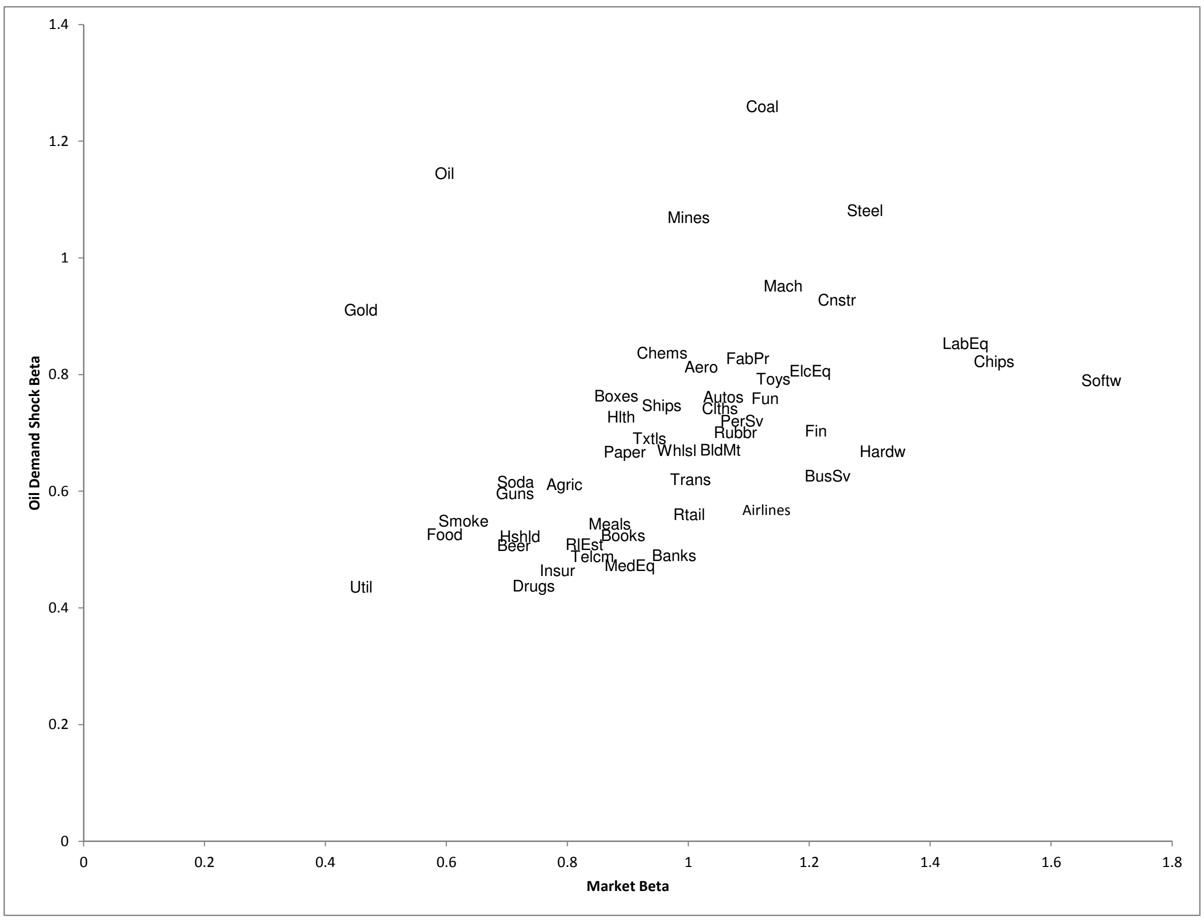

Plots of industry oil use and stock market beta against oil demand shock beta. the X-axis represents the industry beta with aggregate stock market.

the squeeze it puts on household consumers of oil.

Figure 13 plots the betas of each countries' stock index with respect to an oil demand shock against the beta of each country with respect to a world stock market return. The pattern here is similar to the one seen in industries, with the world market beta being an excellent predictor of the aggregate demand beta.

5 Model Extension

In this section the Model is extended to include a monopolist producer operating along side a competitive sector of production. In the presence of the monopolist, the spot price of oil is

$$P_t = A_t(O_t + O_t^M)^{-\frac{1}{\alpha}} \tag{14}$$

Where O_t^M is the oil production of the Monopolist and O_t is the output of the competitive firms. The competitive sector is populated by firms identical to those in the Basic Model. The monopolist produces with the same technology as the competitive producers

$$O_t^M = Z_t^M(F_t^M)^{\nu}(W_t^M)^{1-\nu} \tag{15}$$

Where the log of Z_t^M evolves according to

$$z_{t+1}^M = z_0^M + \rho_z^M(z_t^M - z_0^M) + e_{z,t+1}^M \tag{16}$$

Monopolist oil reserves evolve in the same way as competitive producer returns

$$W_{t+1}^M = W_t^M(1-\delta) + I_t^M - dO_t^M \tag{17}$$

The monopolist profits are given by

$$\Pi_t^M = P_tO_t^M - c_F^M F_t^M - I_t^M - \Phi(F_t^M, F_{t-1}^M, I_t^M, W_t^M) \tag{18}$$

Figure 12: Country Oil Supply Shock Betas

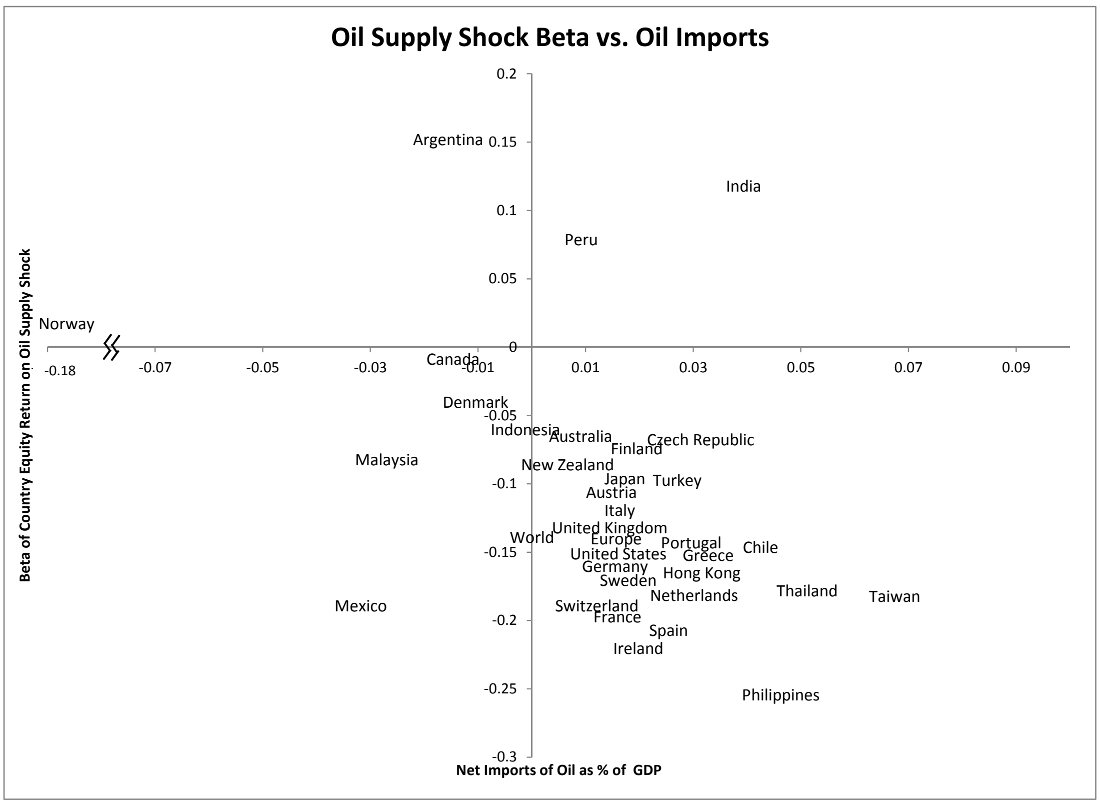

Plots of country scaled oil imports and country equity index world market beta against oil supply shock beta. In panel (a) the X-axis represents the average ratio of oil imports to GDP over the sample period. Data from the EIA and the IMF. In panel (b) the X-axis represents the countries' equity beta with aggregate stock market. Country indices are from Global Financial Data.

Figure 13: Country Oil Demand Shock Betas

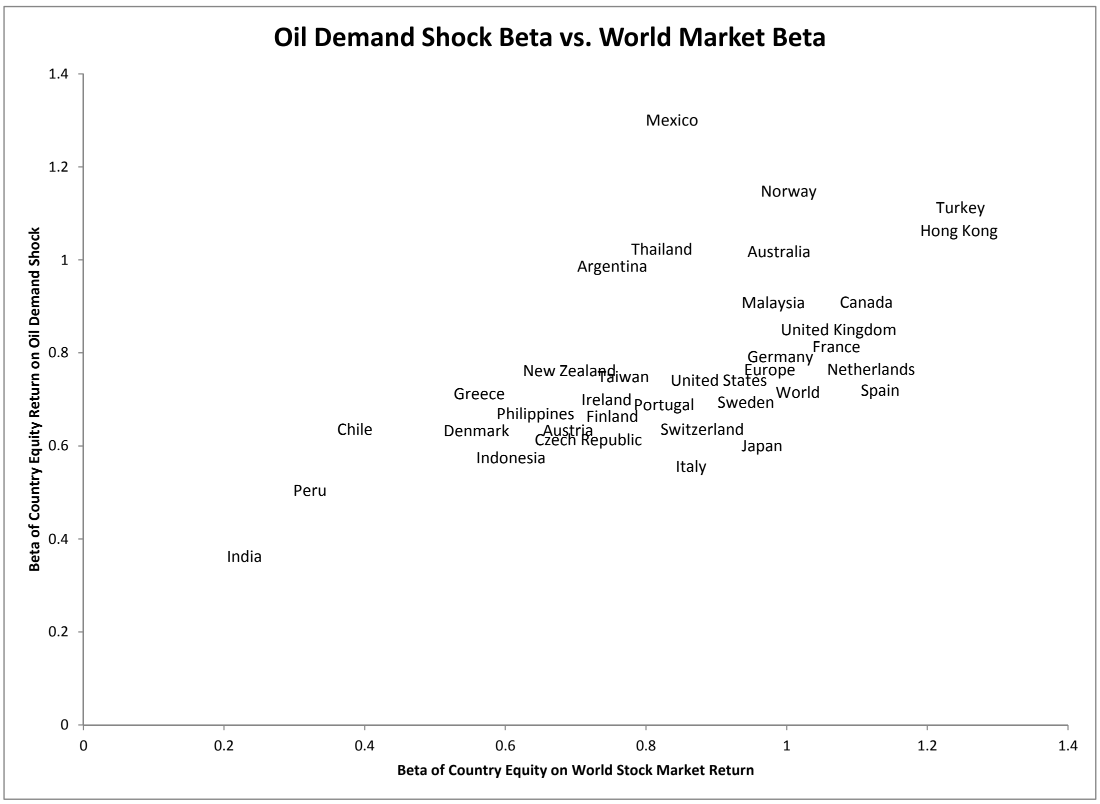

Plots of country scaled oil imports and country equity index world market beta against oil demand shock beta. In panel (a) the X-axis represents the average ratio of oil imports to GDP over the sample period. Data from the EIA and the IMF. In panel (b) the X-axis represents the countries' equity beta with aggregate stock market. Country indices are from Global Financial Data.

Table 9: Model Parameters w/ Monopolist

Parameter	Description	Value
ν	Share of Flow Input in Production	0.6
α	Elasticity of Demand	0.5
σ_a	Volatility of Demand Shock	0.15
σ_z	Volatility of Productivity Shock	0.065
ρ_a	Persistence of Demand Shock	0.95
ρ_z	Persistence of Productivity Shock	0.95
a_0	Mean of Demand Shock	0
z_0	Mean of Productivity Shock	0
c_F	Cost of Flow Input	2.5
a_F	Flow Input Adjustment Cost	3
a_W	Oil Well Adjustment Cost	15
d	Depletion Rate (Benchmark)	1
δ	Deprecation of Oil Wells	0.01
λ_a	Price of Demand Risk	2.0
λ_z	Price of Supply Risk	0.3
β	Discount Rate	0.995

where Φ has the same form and parameterization as in the basic model.

The monopolist takes into account the effect of its production on prices and chooses levels of the flow input and investment to maximize profits. The model is solved numerically for the parameters given in Table 9. The monopolist has a cost advantage in producing the flow input, so that $c_F^M < c_F$. This seems a plausible description of an OPEC producer, and is also necessary for the level of monopolist production relative to the production of the economy to be roughly equivalent to ratio to OPEC production to the rest of the world, which is roughly 45% in the data and 46% in the model.

Impulse response functions for the three shocks in the model are shown in Figure 14. The first two panels show that the shocks to the monopolist's productivity are of minor importance to both the price of oil and the return of the competitive producers.

The shocks to the competitive producers' productivity have a direct impact on prices, but as in the basic model do not affect their stock returns. The model also predicts a strong future rise in monopolist production in response to a demand shock. Perhaps most strikingly, the model implies that shocks to the productivity of the monopolist are very quickly compensated for by an increase in the flow input for Monopolist production, creating only temporary impacts on price. This temporary nature of the shocks means they also have very little impact on competitive oil producer returns. Since neither productivity shock effects the returns to competitive oil producers, they remain an effective control for demand shocks.

In order to understand the effects of the different shocks in the context of the extended model, Figure 15 graphs the response from estimated VARs for OPEC and non-OPEC levels of oil production. Though the production data is very noisy, the patterns in the data are broadly consistent with the model. Supply shocks result in immediate drops in production, which are more persistent for the non-OPEC (presumably competitive) producers, while demand shocks generate a future production increase from OPEC producers.

6 Alternative Specifications of Prices and a Tradeable Factor

In this section different alternatives to the WTI and the Datastream World Integrated Oil and Gas Producer index are used as robustness tests. Additionally, since one of the alternatives is the return to the nearest month NYMEX futures contract, it allows for a tradeable real-time strategy to replicate the constructed supply shock. Though the period for this supply shock (1987 to 2012) is too short for a meaningful investigation of risk premia, it is nevertheless instructive to observe the performance of this claim of the sample.

Figure 14: Model Impulse Response Functions w/ Monopolist

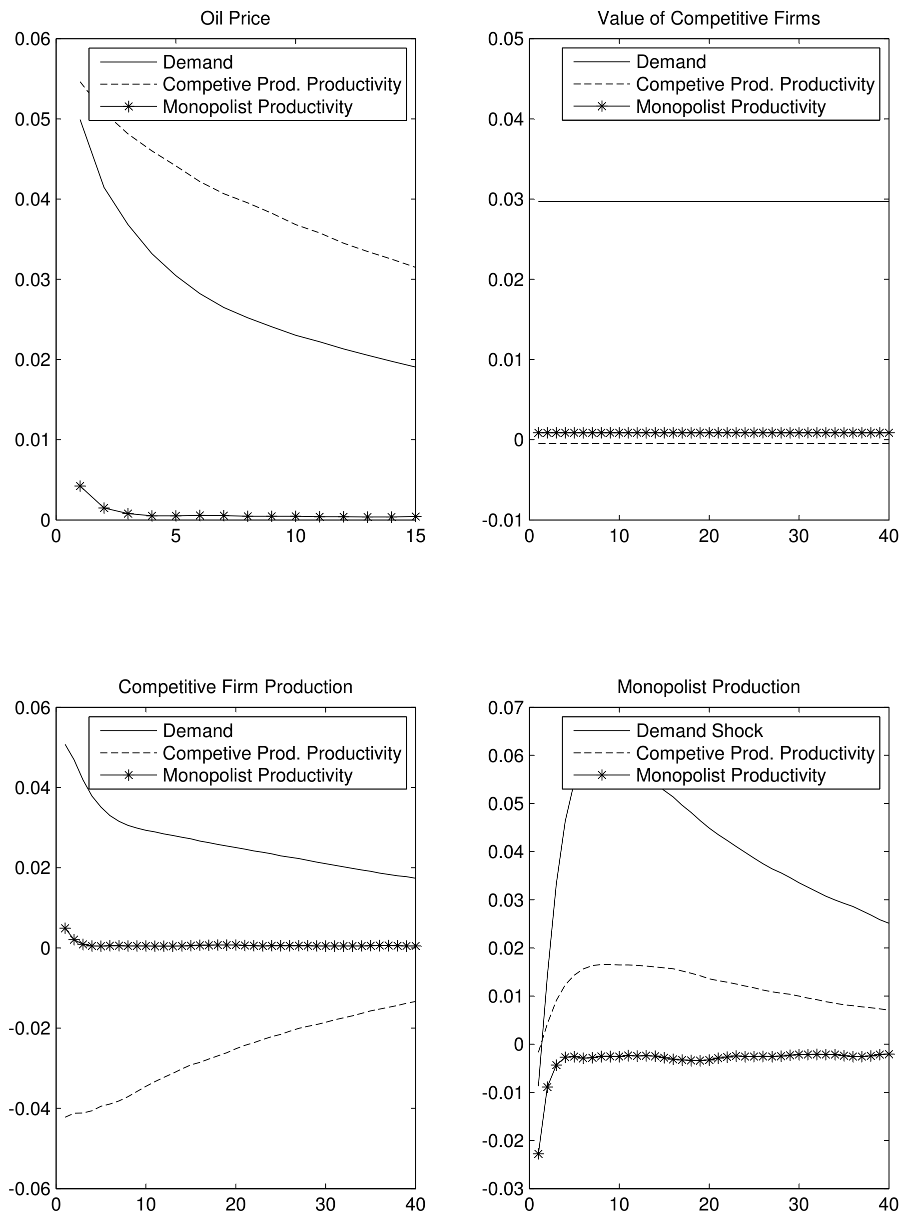

Plots for model impulse responses to demand shocks $e_{a,t}$, and negative productivity shocks to competitive producers $e_{z,t}$, and monopolist $e_{z,t}^{M}$.

Figure 15: OPEC and Non-OPEC Production Responses to Oil Shocks

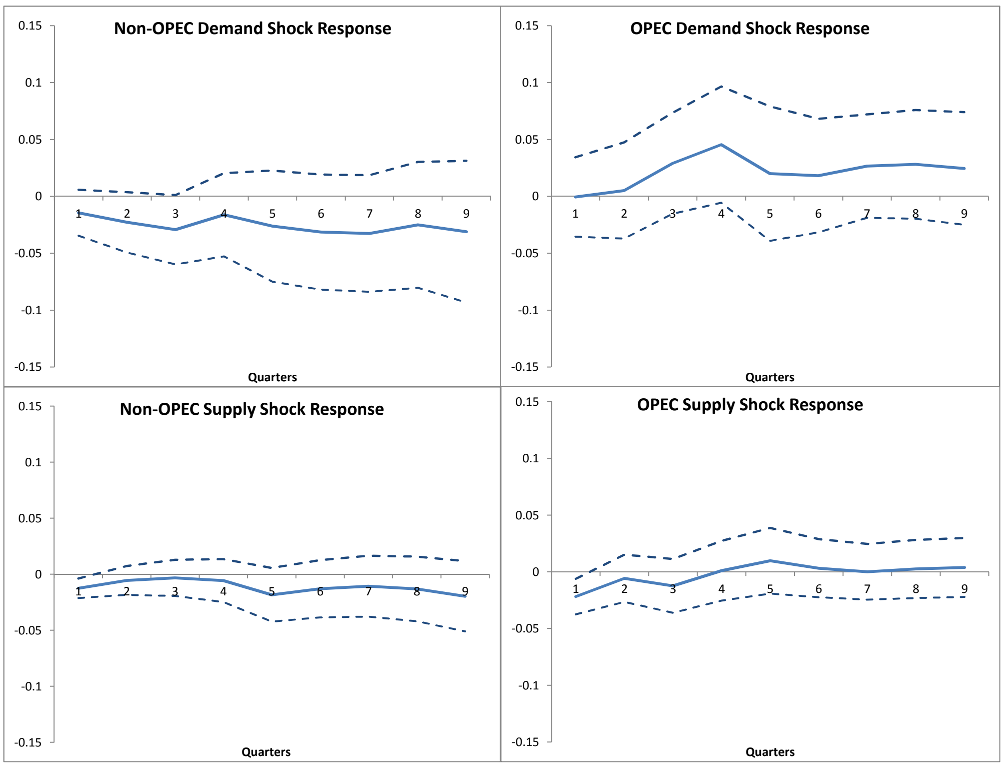

Impulse response for VARs with oil shocks as exogenous variables. Data is monthly from the EIA aggregated to quarterly frequency.

6.1 Alternate Specifications

Table 10 shows the results for univariate regressions of the aggregate U.S. market return against the constructed oil shocks when different choices of oil producer indices and oil price measures are used. For some of the oil price measures the data availability limits the sample so for comparison the benchmark shocks are shown over the same sample. As is clear from the table, the results are robust to these choices, and notably are unaffected by controlling for measures of aggregate price levels, which is done using both the CPI excluding energy and an index of the dollar against a basket of major currencies, available from the St. Louis Fed. Also, as previously mentioned, using the Brent index in the period post crisis yields a very strong negative correlation for stock prices, and although the sample is short this is an interesting development that warrants further observation.

6.2 Performance of Tradeable Strategies

Finally, since the shocks are constructed using rolling regressions and the return to the one month NYMEX future is a tradeable claim, the supply shock constructed using this return represents the return on a tradeable strategy. Namely going long the futures and short a proportional position in the oil producer index. Figure 16 plots the cumulative excess returns from these strategies. Over the 24-year sample, the annualized return on the supply strategy is roughly 2%, which is insignificant. The annualized demand strategy yields roughly 4% annualized returns, which due to the lower volatility is statistically significant, unsurprising since it is very similar to simply taking a long position in the market. Neither of the returns generate CAPM alphas over the period.[16]

[16] A better measure of oil price risk is the slope of the futures curve. See Ready (2010) for a detailed discussion of risk over the sample period.

Table 10: Alternate Price and Index Specifications

Producer Index	**Oil Price Measure**	**Univariate Regressions**			
		Supply Shock		Demand Shock	
		Beta	R^2	Beta	R^2
	Period: 1983-2012				
Benchmark	Benchmark	-0.134**	0.063	0.643**	0.378
		(0.039)		(0.064)	
Benchmark	WTI / Adjusted for Non - Energy CPI	-0.133**	0.063	0.915**	0.315
		(0.039)		(0.101)	
Benchmark	WTI / Adjusted for changes in $	-0.128**	0.062	0.561**	0.377
		(0.042)		(0.056)	
FF Oil Industry	Benchmark	-0.125**	0.057	0.668**	0.378
		(0.041)		(0.063)	
DS World E&P	Benchmark	-0.127**	0.057	0.687**	0.386
		(0.038)		(0.065)	
	1989-2012				
Benchmark	Benchmark	-0.109**	0.042	0.623**	0.270
		(0.039)		(0.058)	
Benchmark	Near Month NYMEX Futures Returns	-0.089*	0.040	0.587**	0.362
		(0.035)		(0.055)	
Benchmark	Brent Oil	-0.103**	0.036	0.625**	0.352
		(0.039)		(0.060)	
	2009-2012				
Benchmark	Benchmark	-0.136	0.018	0.587**	0.598
		(0.169)		(0.093)	
Benchmark	Brent Oil	-0.314*	0.119	0.566**	0.602
		(0.131)		(0.093)	

Univariate regressions of aggregate U.S. stock market returns against supply and demand shocks constructed using different series for oil producer returns and oil prices. Brent crude data and NYMEX WTI futures are from bloomberg. For the $ adjustment to prices the index of major currencies from the St. Louis Fed is used. The CPI is the CPI excluding energy available from the BEA.

Figure 16: Cumulative Returns to Tradeable Strategies

Tradeable strategies are constructed by going long proportional position in oil producers (the demand shock), and long a position in oil futures and short the demand shock (the supply shock). Cumulative returns are shown in excess of the risk free rate (the rate on the one month treasury).

7 Conclusion

This paper presents a new, simple, method for identifying sources of oil supply shocks as the increase in oil prices in excess of the change predicted by the contemporaneous return on oil producers. Using this method, oil supply shocks are shown to have a highly significant impact on U.S. and world stock prices, in contrast to the very small correlations observed when using aggregate changes in oil prices.

Furthermore, this impact is highest on the international level among countries with a high dependence on oil imports, and on the domestic level among firms that depend on consumer expenditure rather than those which rely on oil as an input. These findings provide insight into the way oil price effect the world economy, and suggest that the important effect of oil price shocks may be pain at the pump for consumers rather than higher prices for oil using firms.

The construction of the shocks also provides an new technique for researchers studying the effects of changes in oil prices, particularly those interested in studying impacts at frequencies higher than quarterly data, or over short sample periods.

References

Baker, Steven D., and Bryan R. Routledge, 2011, The price of oil risk, Working paper.

Carlson, Murray, Zeigham Khokher, and Sheridan Titman, 2007, Equilibrium exhaustible resource price dynamics, *Journal of Finance* 62, 1663–1703.

Casassus, Jaime, Pierre Collin-Dufresne, and Bryan R. Routledge, 2005, Equilibrium Commodity Prices with Irreversible Investment and Non-Linear Technology, *SSRN eLibrary*.

Cavallo, Michele, and Tao Wu, 2006, Measuring oil-price shocks using market-based information, Discussion paper.

Chen, Nai-Fu, Richard Roll, and Stephen A. Ross, 1986, Economic forces and the stock market, *Journal of business* pp. 383–403.

Ghoddusi, Hamed, 2010, Dynamic investment in extraction capacity of exhaustible resources, *Scottish Journal of Political Economy* 57, 359–373.

Hamilton, James D, 1983, Oil and the macroeconomy since world war ii, *Journal of Political Economy* 91, 228–48.

Hamilton, James D., 2003, What is an oil shock?, *Journal of Econometrics* 113, 363 – 398.

——— , 2008, Understanding crude oil prices, NBER Working Papers 14492 National Bureau of Economic Research, Inc.

——— , and Jing Cynthia Wu, 2012, Effects of index-fund investing on commodity futures prices, Working paper.

Hotelling, Harold, 1931, The economics of exhaustible resources, *The Journal of Political Economy* 39, 137–175.

Huang, Roger, Ronald W. Masulis, and Hans R. Stoll, 1996, Energy shocks and financial markets, *Journal of Futures Markets* 16, 1–27.

Jones, C., and G. Kaul, 1996, Oil and the stock markets, *Journal of Finance* 51, 463–491.

Kilian, Lutz, 2009, Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market, *The American Economic Review* 99, 1053–1069.

——— , and Cheolbeom Park, 2009, The impact of oil price shocks on the u.s. stock market*, *International Economic Review* 50, 1267–1287.

Kogan, Leonid, Dmitry Livdan, and Amir Yaron, 2009, Oil futures prices in a production economy with investment constraints, *The Journal of Finance* 64, pp. 1345–1375.

Ready, Robert C., 2010, Oil prices and long-run risk, Working paper.

Routledge, Bryan R., Duane J. Seppi, and Chester S. Spatt, 2000, Equilibrium forward curves for commodities, *The Journal of Finance* 55, 1297–1338.

Sadorsky, Perry, 1999, Oil price shocks and stock market activity, *Energy Economics* 21, 449–469.

White, Halbert, 1980, A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity, *Econometrica* 48, pp. 817–838.

Williams, J.C., and B.D. Wright, 1991, *Storage and Commodity Markets* (Cambridge University Press: Cambridge, UK).